UNIVERSITY OF TWENTE.

DEEPCASE: Semi-Supervised Contextual Analysis of Security Events

Thijs van Ede, Hojjat Aghakhani, Noah Spahn, Riccardo Bortolameotti, Marco Cova, Andrea Continella, Maarten van Steen, Andreas Peter, Christopher Kruegel and Giovanni Vigna

Contact: t.s.vanede@utwente.nl

REAQTA lastline vmware

DeepCASE

- Goal: Reduce the workload of Security Operators in SOCs

How does a SOC operate?

- Goal: Reduce the workload of Security Operators in SOCs
- Collect activity from devices

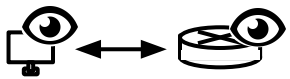

Collect activity from network and hosts:

- `PROGRAM $X OPENED FILE $Y`
- `$IP_A ESTABLISHED A CONNECTION TO $IP_B`

How does a SOC operate?

- Goal: Reduce the workload of Security Operators in SOCs
- Collect activity from devices
- Monitor activity for suspicious patterns

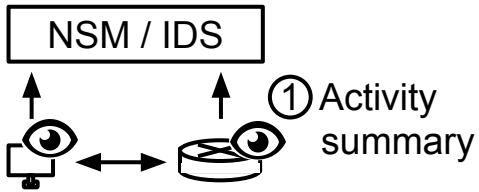

How does a SOC operate?

- Goal: Reduce the workload of Security Operators in SOCs
- Collect activity from devices
- Monitor activity for suspicious patterns
- Send events to SOC

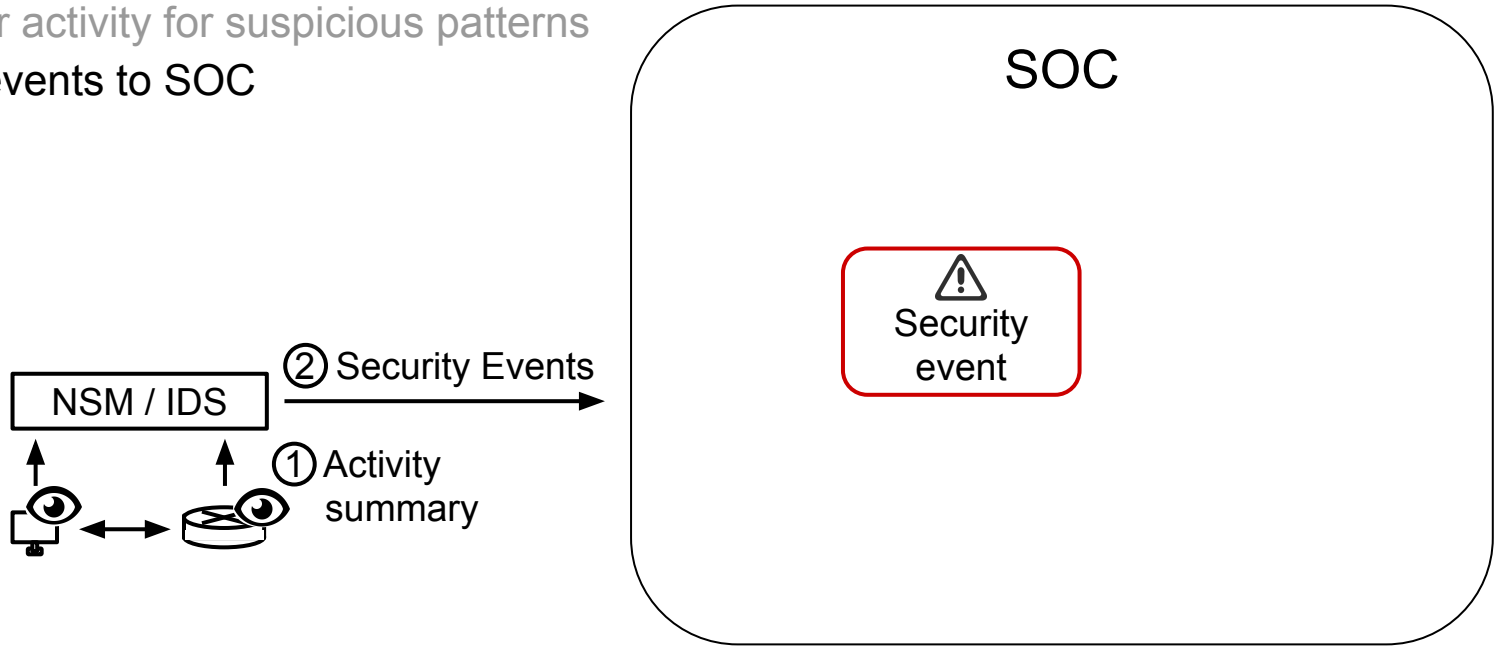

How does a SOC operate?

- Goal: Reduce the workload of Security Operators in SOCs
- Collect activity from devices
- Monitor activity for suspicious patterns
- Send events to SOC, where they get:
  - Triaged

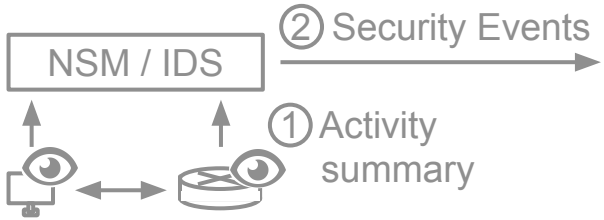

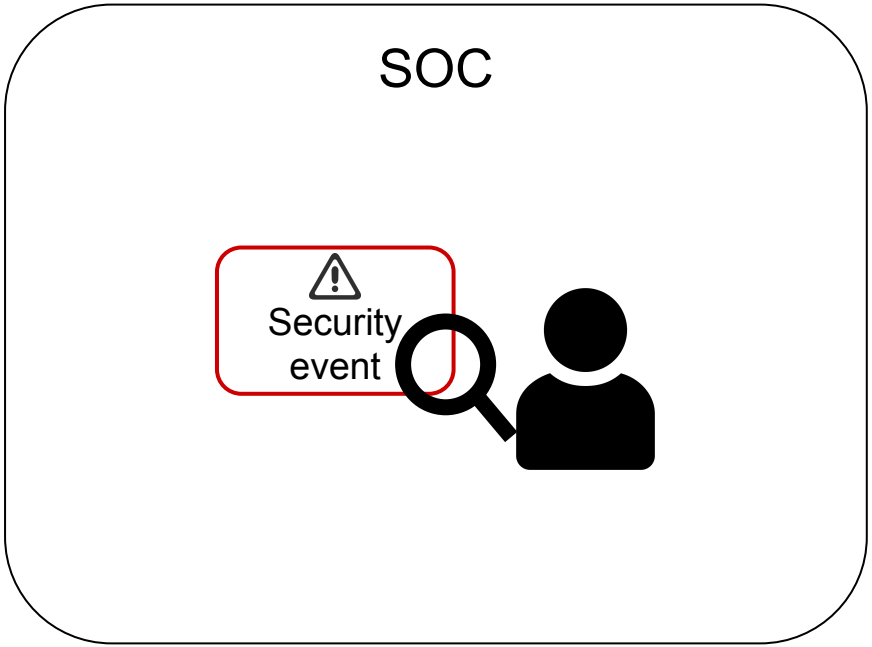

How does a SOC operate?

- Goal: Reduce the workload of Security Operators in SOCs
- Collect activity from devices
- Monitor activity for suspicious patterns
- Send events to SOC, where they get:
  - Triaged
  - Discarded

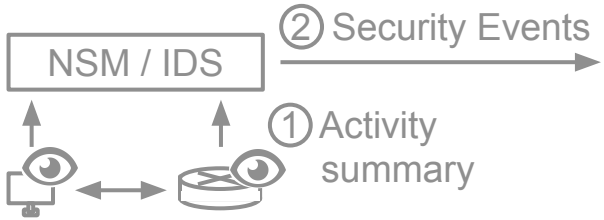

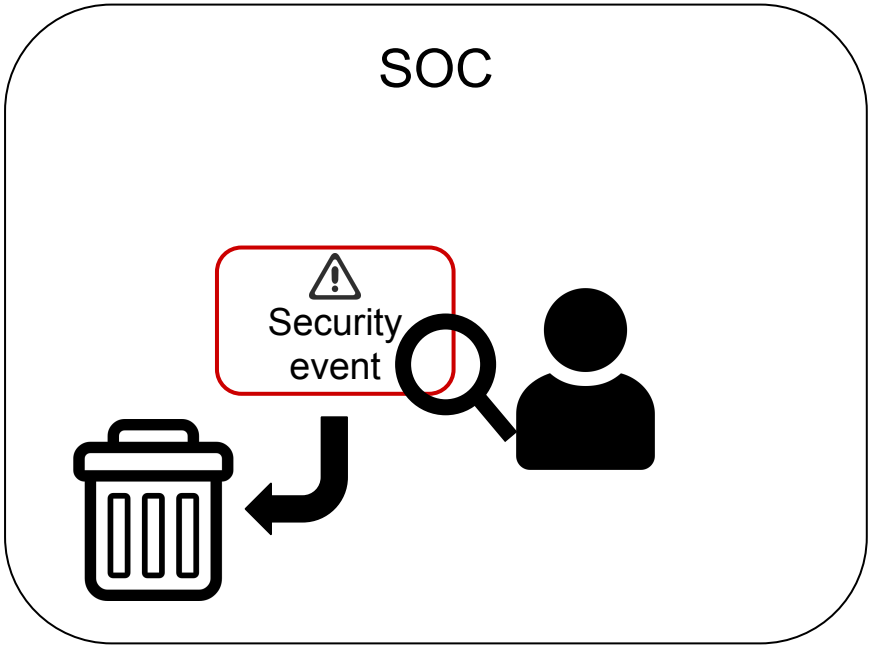

How does a SOC operate?

- Goal: Reduce the workload of Security Operators in SOCs
- Collect activity from devices
- Monitor activity for suspicious patterns
- Send events to SOC, where they get:
  - Triaged
  - Discarded
  - Escalated

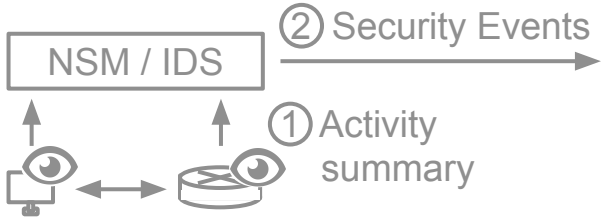

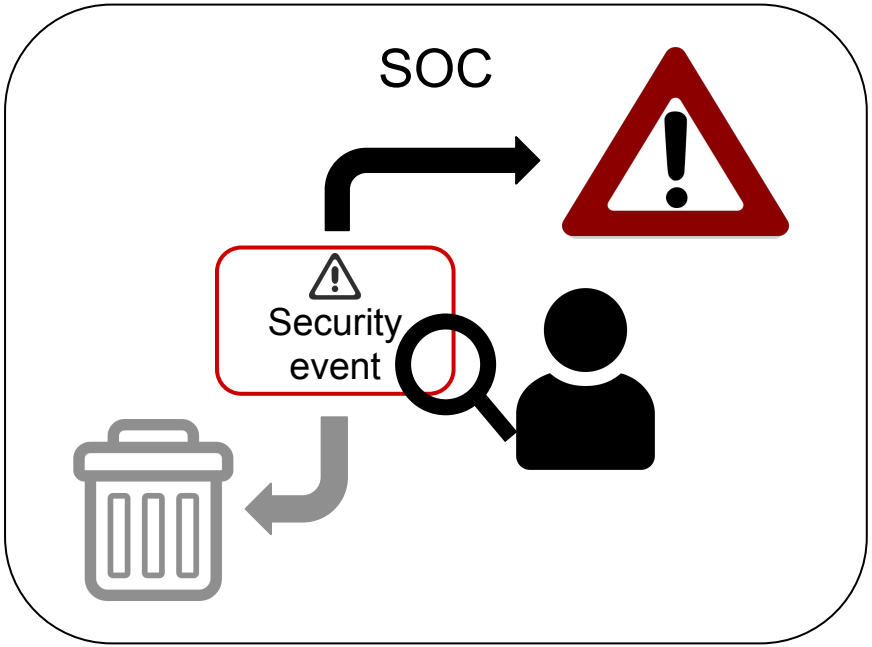

How does a SOC operate?

- Goal: Reduce the workload of Security Operators in SOCs
- Collect activity from devices
- Monitor activity f
- Send events to S
  - Triaged
  - Discarded
  - Escalated

Can we **automate** security event **triaging** to **reduce** the workload of security operators?

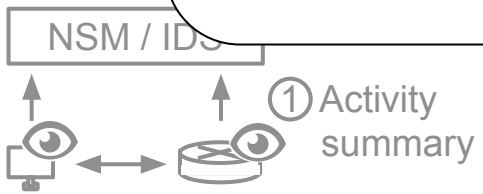

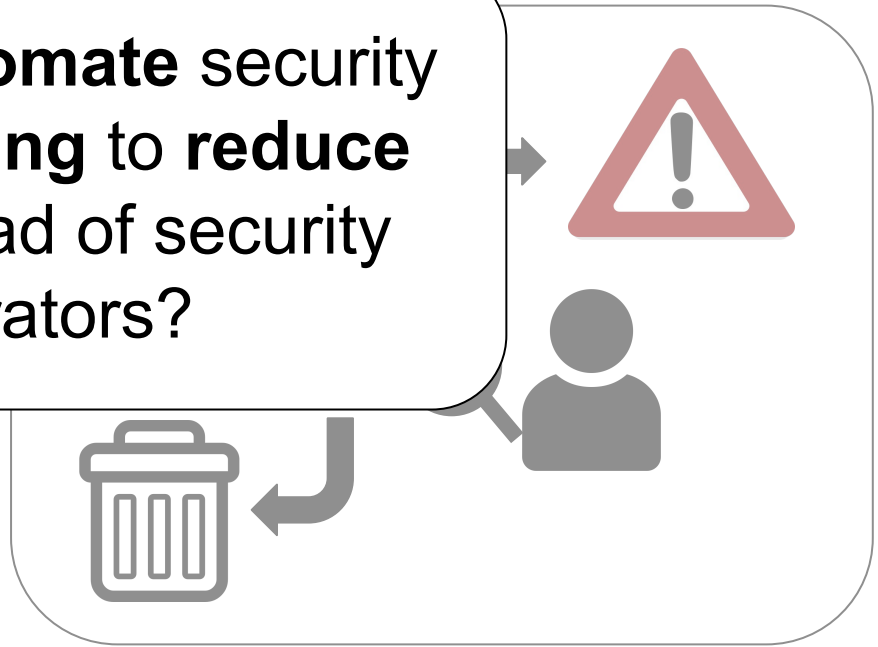

Intuition

- An individual security event may not say much about a potential threat to the system

Intuition

- An individual security event may not say much about a potential threat to the system, but looking at **contextual events** tells us a lot.

Intuition

- An individual security event may not say much about a potential threat to the system, but looking at **contextual events** tells us a lot.

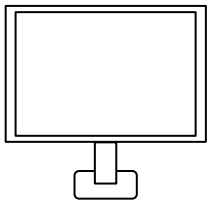

Intuition

- An individual security event may not say much about a potential threat to the system, but looking at **contextual events** tells us a lot.

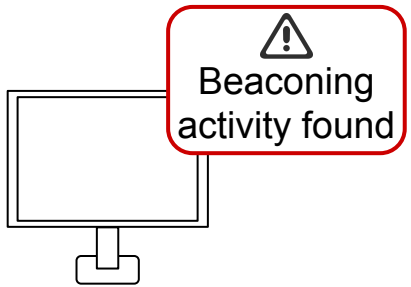

Intuition

- An individual security event may not say much about a potential threat to the system, but looking at **contextual events** tells us a lot.

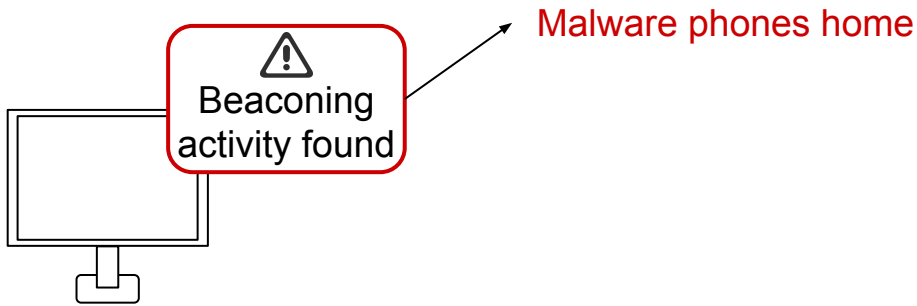

Intuition

- An individual security event may not say much about a potential threat to the system, but looking at **contextual events** tells us a lot.

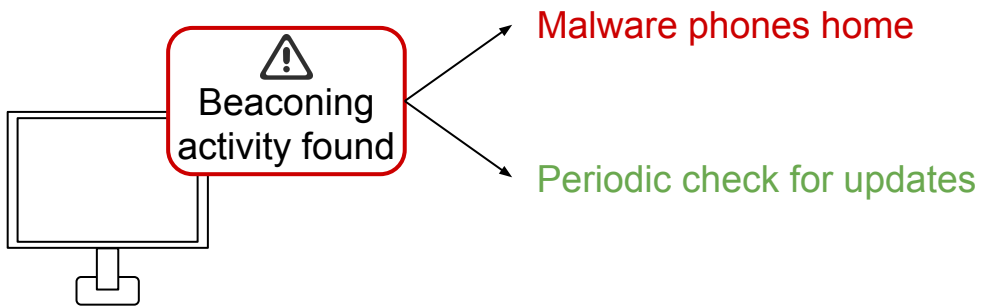

Intuition

- An individual security event may not say much about a potential threat to the system, but looking at **contex**

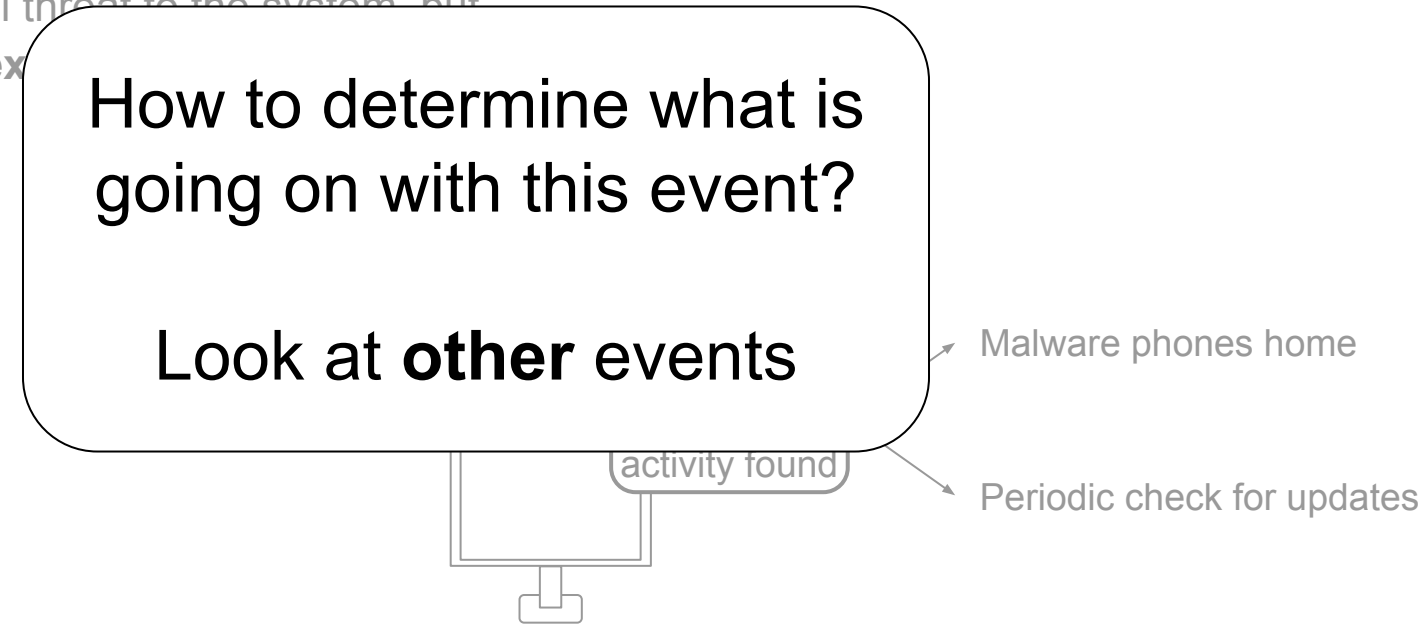

Intuition

- An individual security event may not say much about a potential threat to the system, but looking at **contextual events** tells us a lot.

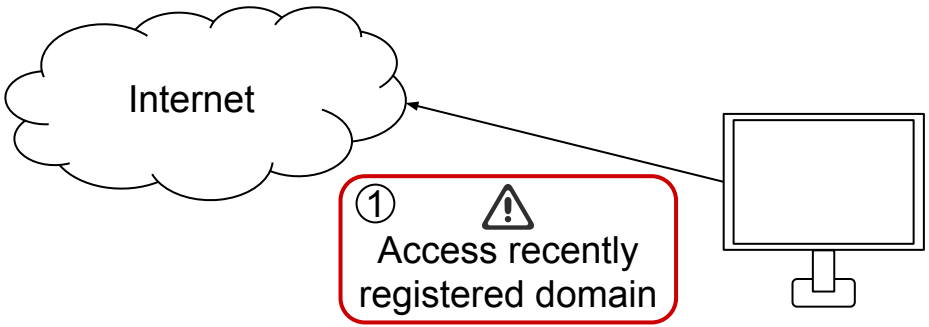

Intuition

- An individual security event may not say much about a potential threat to the system, but looking at **contextual events** tells us a lot.

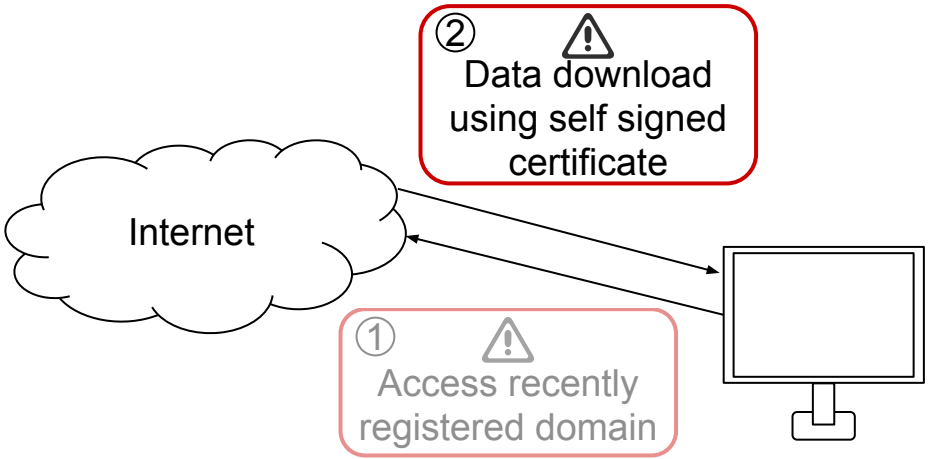

Intuition

- An individual security event may not say much about a potential threat to the system, but looking at **contextual events** tells us a lot.

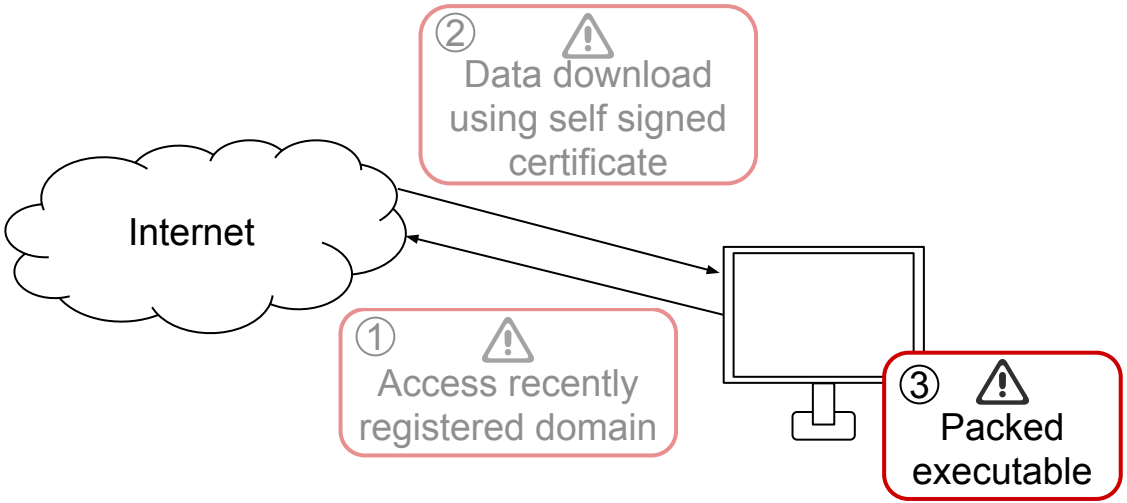

Intuition

- An individual security event may not say much about a potential threat to the system, but looking at **contextual events** tells us a lot.

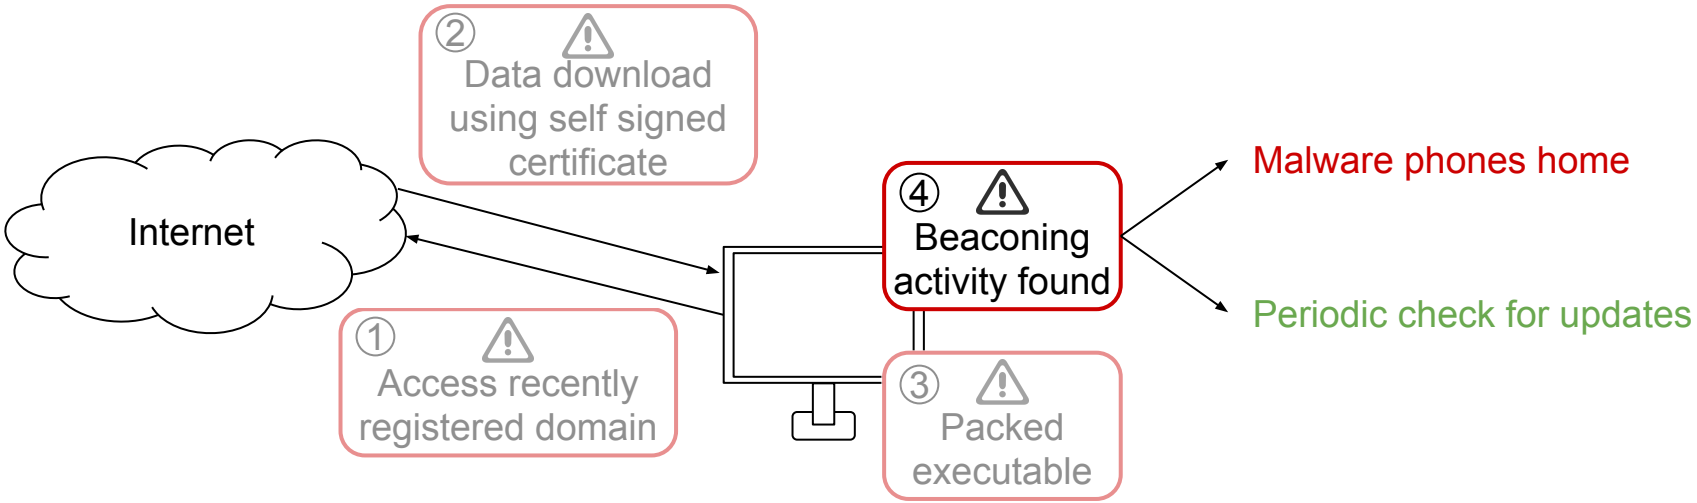

Intuition

- An individual security event may not say much about a potential threat to the system, but looking at **contextual events** tells us a lot.

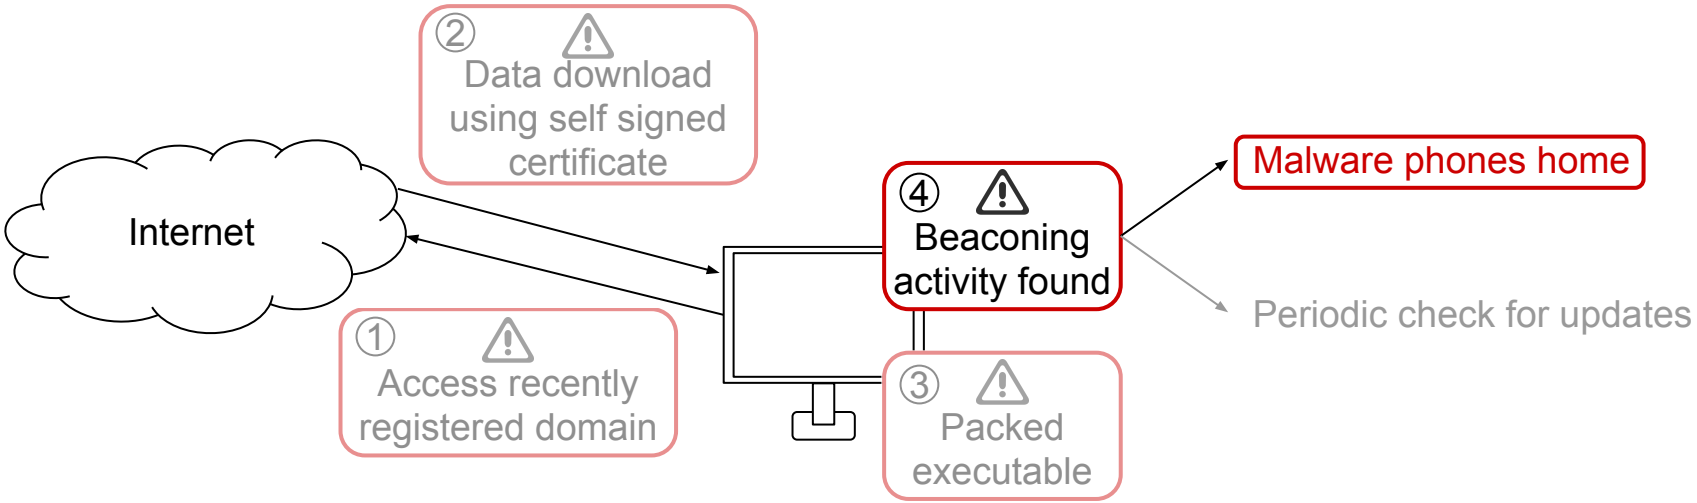

DEEPCASE - Overview

- Idea:
  - Ask security operators to triage events based on their context

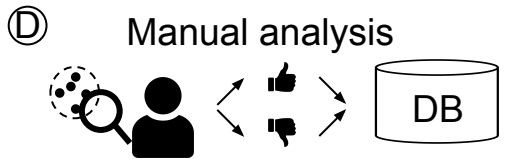

DEEPCASE - Overview

- Idea:
  - Ask security operators to triage events based on their context
  - Automatically escalate / discard similar context + event to **reduce** operator **workload**

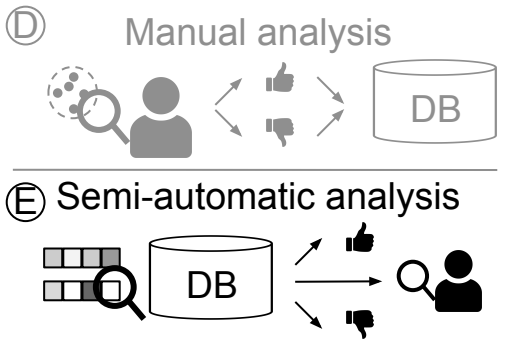

DEEPCASE - Overview

- Idea:
  - Ask security operators to triage events based on their context
  - Automatically escalate / discard similar context + event to **reduce** operator **workload**
  - We need some definition of “similar”

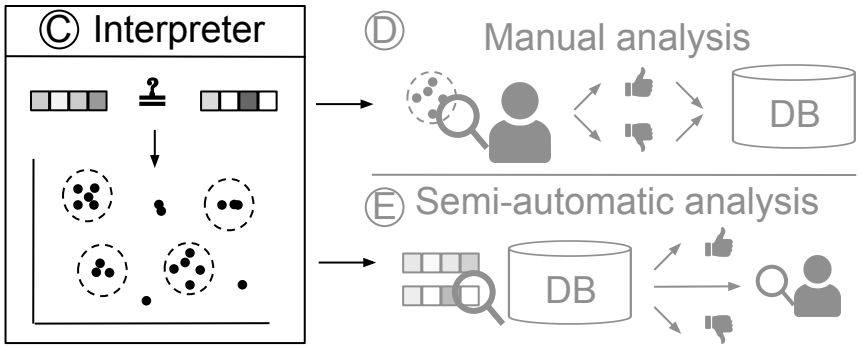

DEEPCASE - Overview

- Idea:
  - Ask security operators to triage events based on their context
  - Automatically escalate / discard similar context + event to **reduce** operator **workload**
  - We need some definition of “similar”
  - We need to deal with irrelevant contextual events

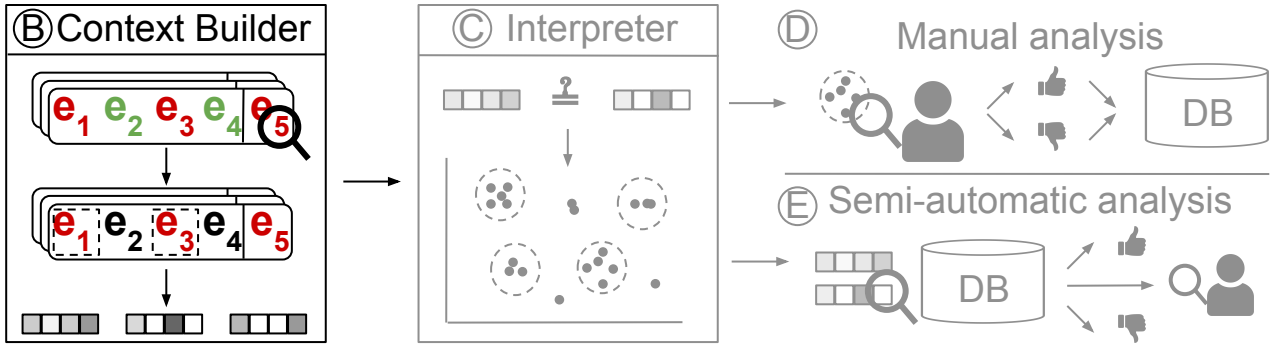

DEEPCASE - Overview

- Idea:
  - Ask security operators to triage events based on their context
  - Automatically escalate / discard similar context + event to **reduce** operator **workload**
  - We need some definition of “similar”
  - We need to deal with irrelevant contextual events
  - We need to collect the events

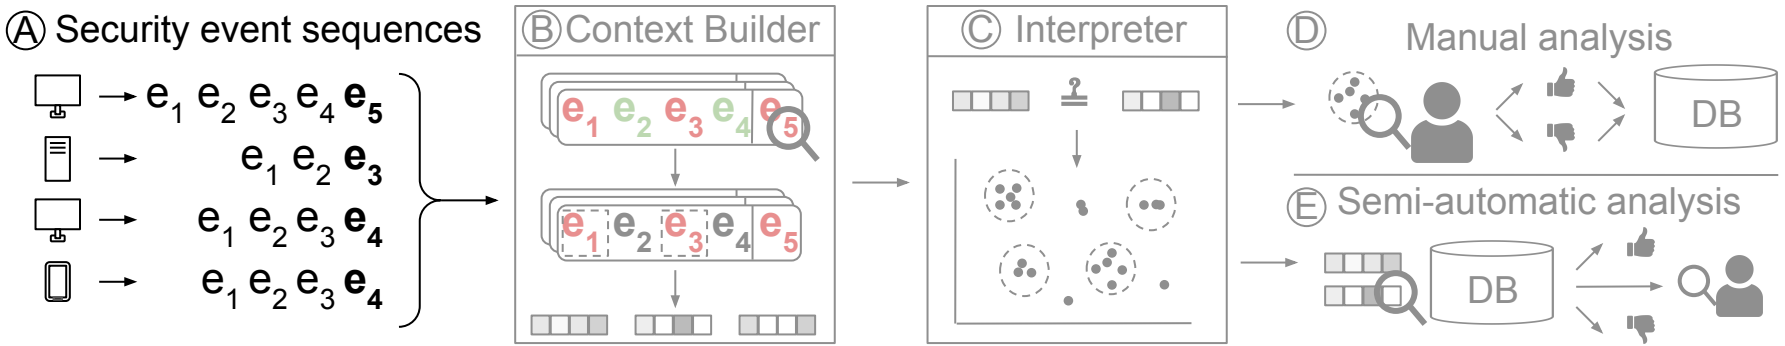

DEEPCASE - Security event sequences

- Collect events per device, sorted by time

Ⓐ Security event sequences

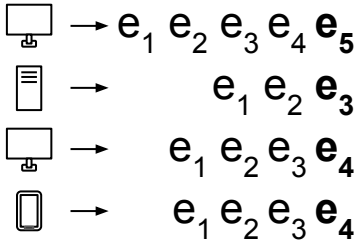

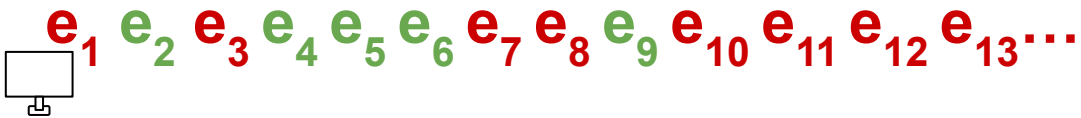

DEEPCASE - Security event sequences

- Collect events per device, sorted by time
- Create a sliding window (sequence) over events

Ⓐ Security event sequences

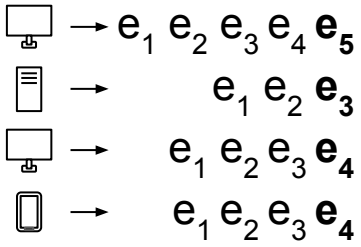

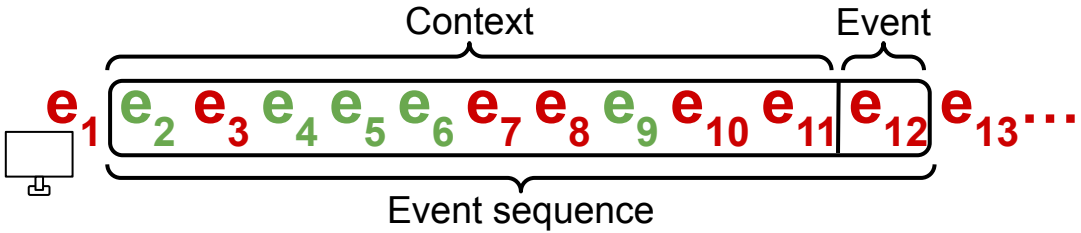

DEEPCASE - Security event sequences

- Collect events per device, sorted by time
- Create a sliding window (sequence) over events:
  - Context length (10 events)
  - Time (1 day)

Ⓐ Security event sequences

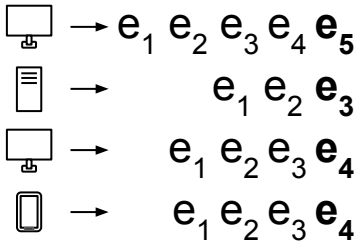

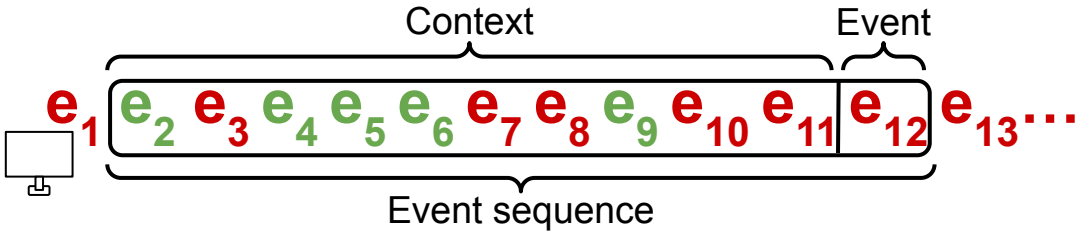

DEEPCASE - Context Builder

- Deal with irrelevant contextual events

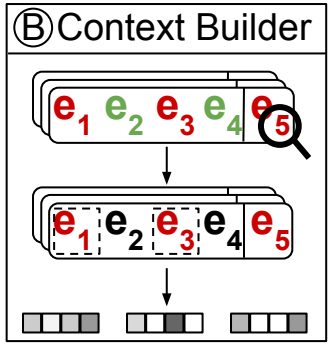

DEEPCASE - Context Builder

- Deal with irrelevant contextual events
- Idea: Train an algorithm to predict an event from the preceding, **contextual** events

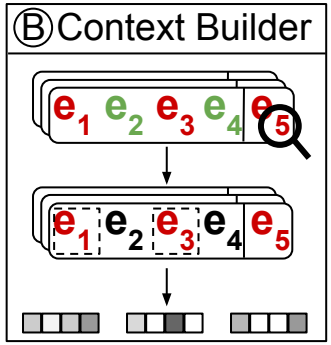

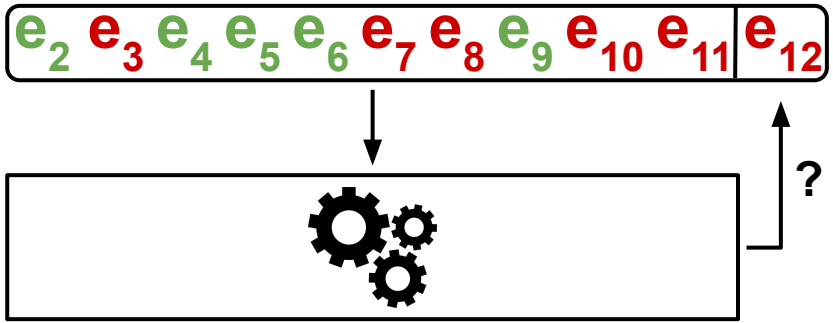

DEEPCASE - Context Builder

- Deal with irrelevant contextual events
- Idea: Train an algorithm to predict an event from the preceding, **contextual** events, then look at the events on which the algorithm **focused** during correct prediction.

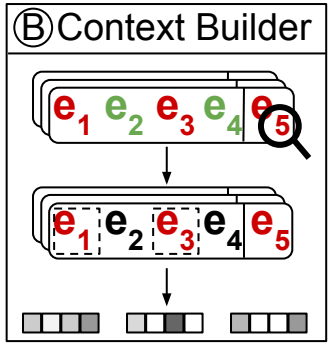

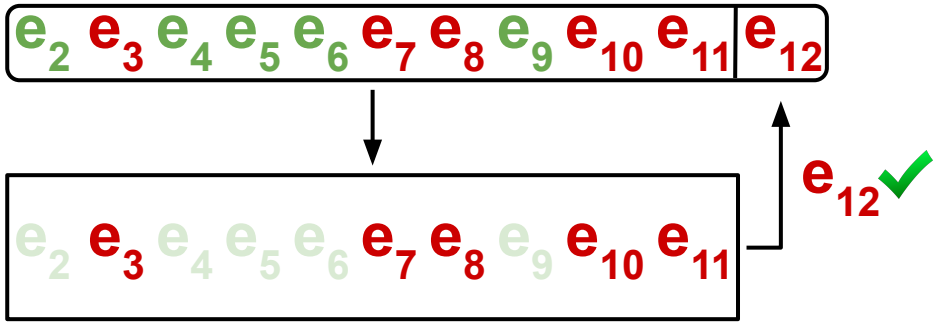

DEEPCASE - Context Builder

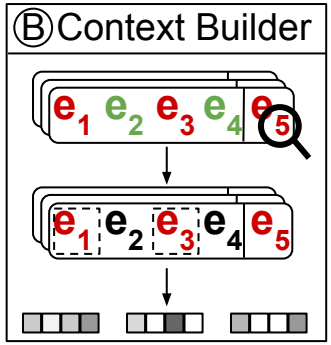

- Deal with irrelevant contextual events
- Idea: Train an algorithm to predict an event from the preceding, **contextual** events, then look at the events on which the algorithm **focused** during correct prediction. These events are likely **correlated** with the prediction.

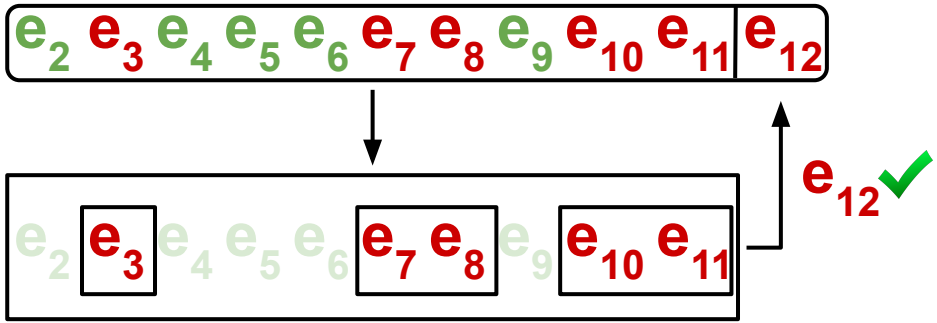

DEEPCASE - Context Builder

- Neural network: Attention-based encoder-decoder model

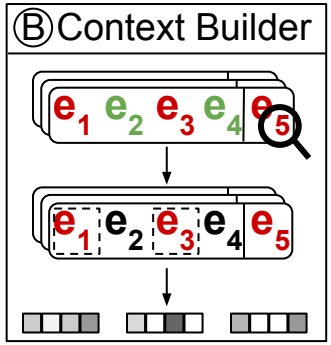

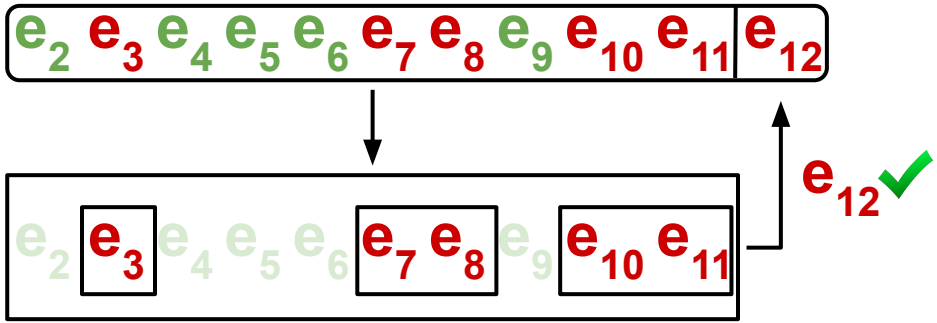

DEEPCASE - Context Builder

- Neural network: Attention-based encoder-decoder model

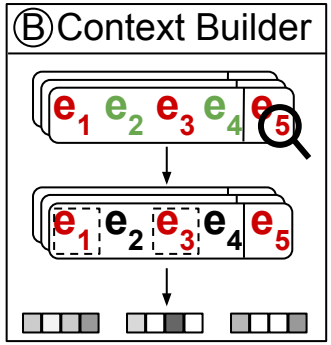

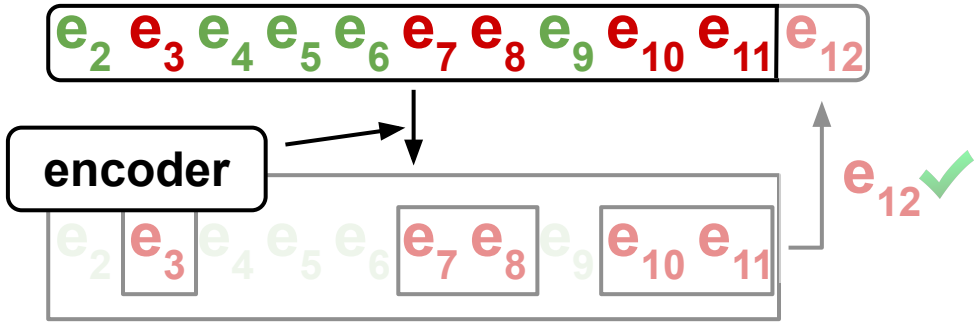

DEEPCASE - Context Builder

- Neural network: Attention-based encoder-decoder model

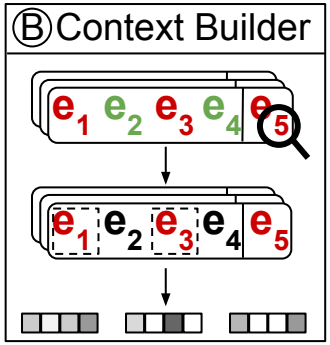

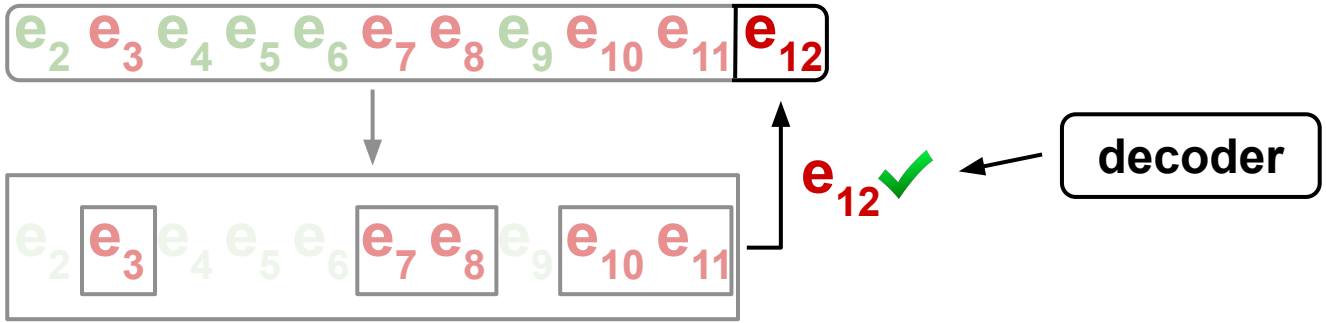

DEEPCASE - Context Builder

- Neural network: Attention-based encoder-decoder model

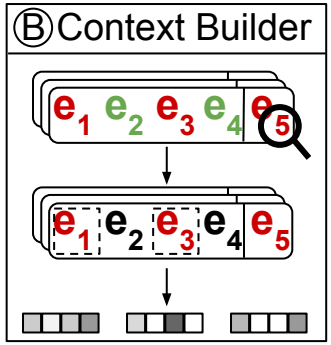

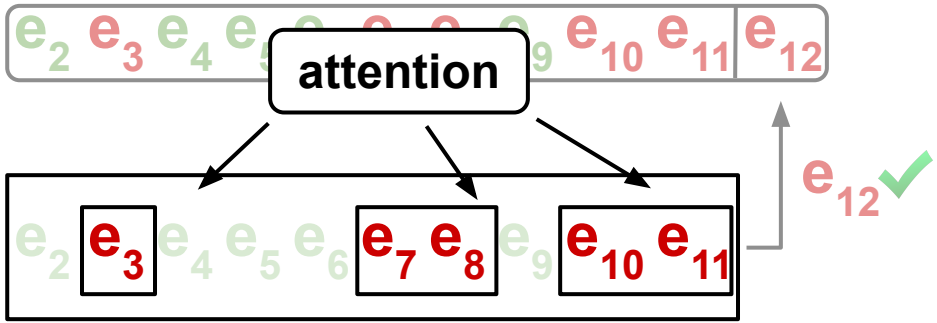

DEEPCASE - Context Builder

- Neural network: Attention-based encoder-decoder model
  - The network **automatically** learns how to distribute attention

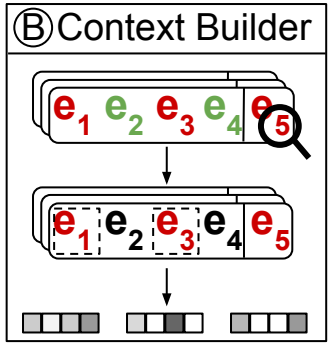

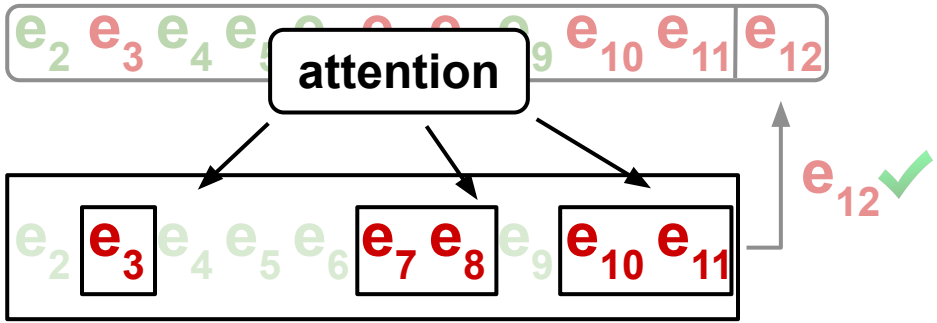

DeepCASE - Context Builder

- Neural network: Attention-based encoder-decoder model
  - The network automatically learns how to distribute attention
- Output: ∑ attention × encoded events

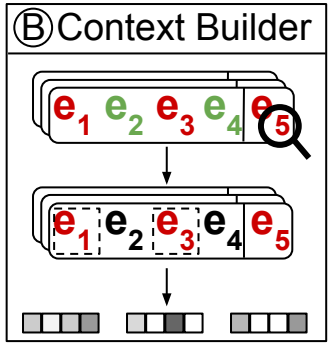

e_2 e_3 e_4 e_5 e_6 e_7 e_8 e_9 e_{10} e_{11} | e_{12}

Event	Type	Attention Weight
e_2	RDP connection	0.01
e_3	Port scan	0.39
e_4	RDP connection	0.02
…	…	…

Sum →

Type	Total Attention Weight
RDP connection	0.03
Port scan	0.39
…	…

DEEPCASE - Interpreter

- Group together alerts from similar event sequences

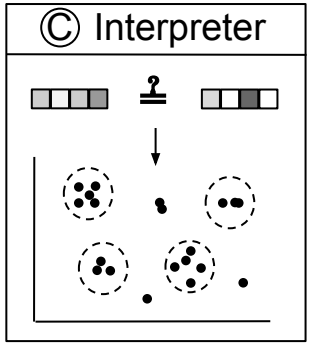

DEEPCASE - Interpreter

- Group together alerts from similar event sequences
- When are event sequences similar?

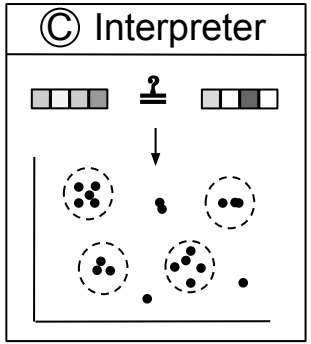

DEEPCASE - Interpreter

- Group together alerts from similar event sequences
- When are event sequences similar?
  - If the attack happens in a similar context

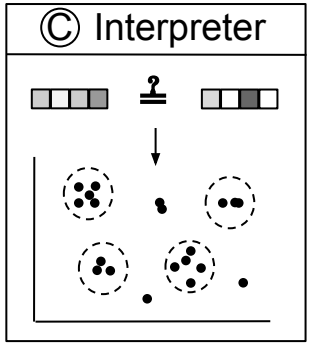

DEEPCASE - Interpreter

- Group together alerts from similar event sequences
- When are event sequences similar?
  - If the attack happens in a similar context
  - Measure through **attention vector**

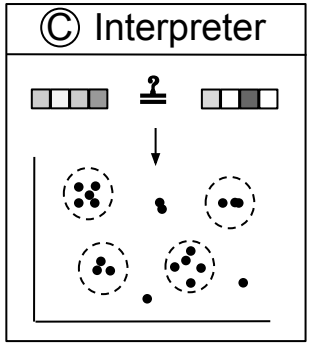

DEEPCASE - Interpreter

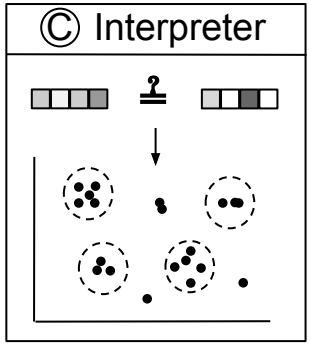

- Group together alerts from similar event sequences
- When are event sequences similar?
  - If the attack happens in a similar context
  - Measure through **attention vector**
  - Compute **manhattan distance**

$$d_1(\mathbf{p}, \mathbf{q}) = \|\mathbf{p} - \mathbf{q}\|_1 = \sum_{i=1}^{n} |p_i - q_i|,$$

DEEPCASE - Interpreter

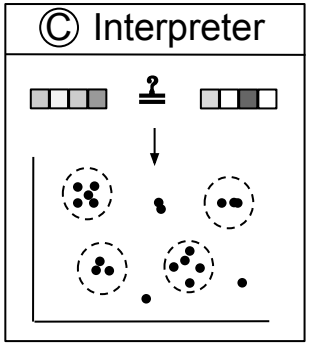

- Group together alerts from similar event sequences
- When are event sequences similar?
  - If the attack happens in a similar context
  - Measure through **attention vector**
  - Compute **manhattan distance**

$$d_1(\mathbf{p}, \mathbf{q}) = \|\mathbf{p} - \mathbf{q}\|_1 = \sum_{i=1}^{n} |p_i - q_i|,$$

- Cluster similar event sequences together

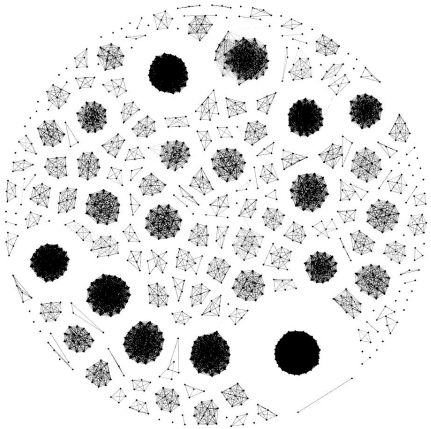

DEEPCASE - Manual Analysis

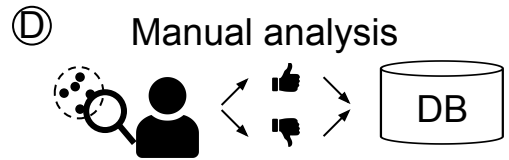

- The Security Operator triages event sequences

DEEPCASE - Manual Analysis

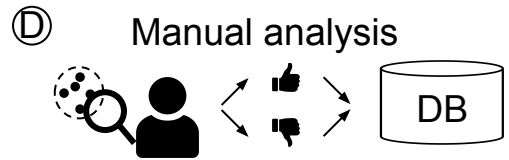

- The Security Operator triages event sequences
- Sample alerts in each cluster for manual analysis

DEEPCASE - Manual Analysis

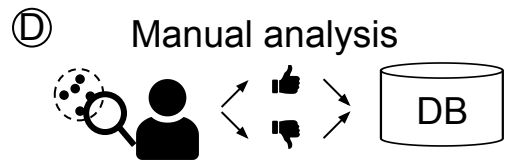

- The Security Operator triages event sequences
- Sample alerts in each cluster for manual analysis
- Not all clustered sequences have the same risk

		# Sequences				
Risk level	Clusters	Total	Average	Min	Max	σ (SD)
INFO	1,115	1.216M	1090.3	5	583.9K	19.2K
LOW	221	41.8K	189.4	5	5,557	612.9
MEDIUM	18	568	31.6	5	235	55.5
HIGH	17	1989	117.0	6	1,107	270.6
ATTACK	33	1391	42.2	5	402	77.1
SUSPICIOUS	238	619.8K	2604.4	5	280.1K	20.2K
Total	1,642	1.881M	1145.7	5	583.9K	17.6K

14.5% of clusters in our analysis contain events of **different** risk levels

DEEPCASE - Manual Analysis

- The Security Operator triages event sequences
- Sample alerts in each cluster for manual analysis
- Not all clustered sequences have the same risk
- Safe approach: always escalate suspicious clusters to security operator

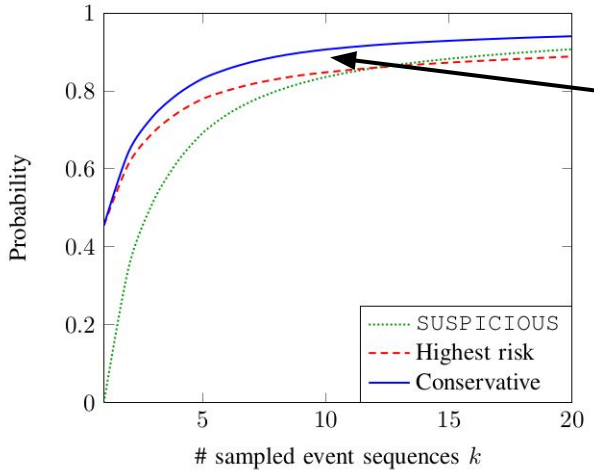

Inspecting only 10 sequences per cluster gave a 84.5% confidence of finding a `SUSPICIOUS` cluster

DEEPCASE - Manual Analysis

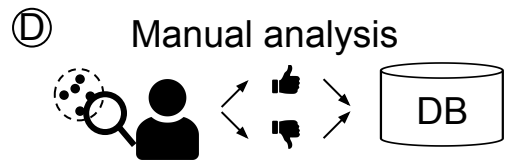

- The Security Operator triages event sequences
- Sample alerts in each cluster for manual analysis
- Not all clustered sequences have the same risk
- Safe approach: always escalate suspicious clusters to security operator
- Inspecting 10 samples per cluster **reduced** triaging workload by **95.39%**

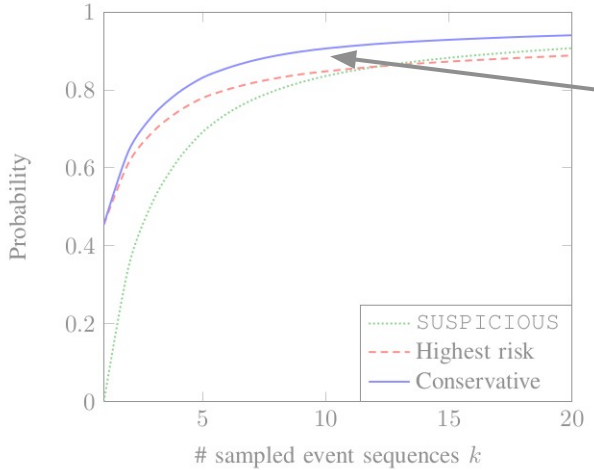

Inspecting only 10 sequences per cluster gave a 84.5% confidence of finding a SUSPICIOUS cluster

DEEPCASE - Semi-automatic Analysis

Ⓔ Semi-automatic analysis

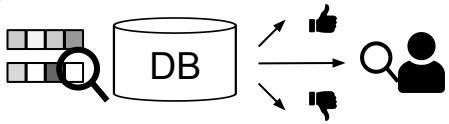

- Automatically escalate / discard “known” alerts

DEEPCASE - Semi-automatic Analysis

Ⓔ Semi-automatic analysis

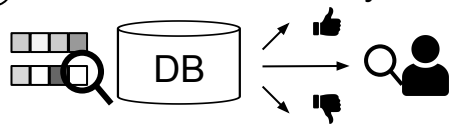

- Automatically escalate / discard “known” alerts
- Manually inspect “new” alerts

DEEPCASE - Semi-automatic Analysis

Ⓔ Semi-automatic analysis

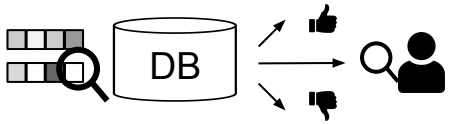

- Automatically escalate / discard “known” alerts
- Manually inspect “new” alerts
- Performance:
  - 86.72% of sequences can be handled automatically

DEEPCASE - Semi-automatic Analysis

Ⓔ Semi-automatic analysis

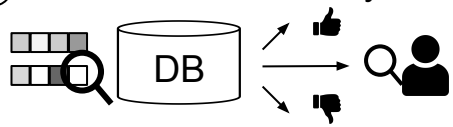

- Automatically escalate / discard “known” alerts
- Manually inspect “new” alerts
- Performance:
  - 86.72% of sequences can be handled automatically
  - Combined with manual inspection of new alerts, workload is reduced by 90.53%

DEEPCASE - Semi-automatic Analysis

Ⓔ Semi-automatic analysis

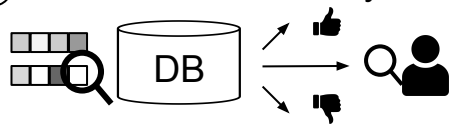

- Automatically escalate / discard “known” alerts
- Manually inspect “new” alerts
- Performance:
  - 86.72% of sequences can be handled automatically
  - Combined with manual inspection of new alerts, workload is reduced by 90.53%
  - We underestimate less than 0.001% of security risks

Conclusion

DEEPCASE **reduces** workload of security operators by analyzing contextual security events

- **Reduces** triaging workload of security operators by 95.39%
- **Automatically** handles 90.53% of events
- Underestimates security risks in **less than** 0.001% of cases

`https://github.com/Thijsvanede/DeepCASE`

Questions?

DEEPCASE **reduces** workload of security operators by analyzing contextual security events

- **Reduces** triaging workload of security operators by 95.39%
- **Automatically** handles 90.53% of events
- Underestimates security risks in **less than** 0.001% of cases

`https://github.com/Thijsvanede/DeepCASE`

Thijs van Ede
t.s.vanede@utwente.nl
@EdeThijs

REAQTA lastline vmware UNIVERSITY OF CALIFORNIA SANTA BARBARA UNIVERSITY OF TWENTE.

Ⓔ Semi-automatic analysis

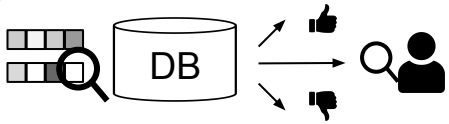

DEEPCASE - Semi-automatic Analysis

- Automatically escalate / discard “known” alerts
- Manually inspect “new” alerts
- Performance:

		Workload reduction				Performance over covered events				
	Method	**Alerts[A]**	**Reduction[B]**	**Coverage[C]**	**Overall[D]**	**Precision**	**Recall**	**F1-score**	**Accuracy**	**Underest.**
Semi-automatic	DEEPCASE	51,800	99.19%	91.27%	**90.53%**	96.39%	91.47%	93.41%	91.47%	**< 0.01%**
	fully-automatic part	N/A	100.00%	86.72%	86.72%	96.39%	91.47%	93.41%	91.47%	**< 0.01%**
	manual part	51,800	83.83%	34.29%	28.74%	N/A	N/A	N/A	N/A	N/A
	Alert throttling (15 min)	3,532,849	49.77%	100.00%	49.77%	98.08%	98.04%	98.04%	98.04%	0.79%
	Alert throttling (1 day)	855,798	87.83%	100.00%	87.83%	97.47%	97.49%	97.49%	97.47%	1.34%
	Rules AlienVault[E]	421,693	83.78%	36.97%	30.97%	99.64%	99.63%	99.63%	99.63%	0.16%
	Rules VMWARE [F]	299,246	89.49%	27.02%	24.18%	100.00%[F]	100.00%[F]	100.00%[F]	100.00%[F]	**0.00%[F]**
	Rules Sigma/Zeek[E]	126,147	92.87%	25.14%	23.35%	99.55%	99.51%	99.52%	99.51%	0.17%