Exchanging Security Events: Which And How Many Alerts Can We Aggregate?

Martin Husák, Milan Čermák, Martin Laštovička, and Jan Vykopal
Institute of Computer Science, Masaryk University
Brno, Czech Republic
{husakm|cermak|lastovicka|vykopal}@ics.muni.cz

Abstract

The exchange of security alerts is a current trend in network security and incident response. Alerts from network intrusion detection systems are shared among organizations so that it is possible to see the "big picture" of current security situation. However, the quality and redundancy of the input data seem to be underrated. We present four use cases of aggregation of the alerts from network intrusion detection systems. Alerts from a sharing platform deployed in the Czech national research and education network were examined in a case study. Volumes of raw and aggregated data are presented and a rule of thumb is proposed: **up to 85 % of alerts can be aggregated**.

Use Case 1 – Duplicates

- Duplicates hold the same information, but are found in multiple copies in the sharing platform.
- They are often caused by sharing the data between more systems. For example, a sensor raises an alert, sends it to the alert sharing platform and a local reporting system, but the local system also sends it to the alert sharing platform.
- Distribution of duplicates per unique alert creates the long tail of alerts duplicated only once.

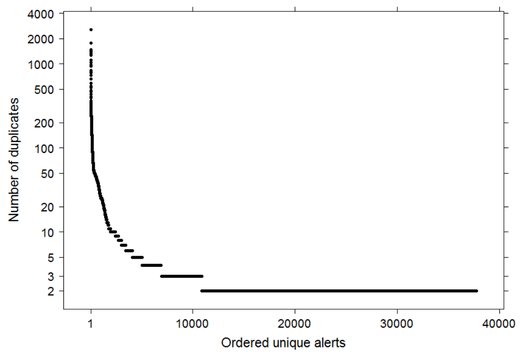

Use Case 2 – Continuing alerts

- Continuing alerts are alerts of the same event that was detected repeatedly.
- For example, an event takes longer time than the processing time window of the detecting sensor and the sensor does not check for its previous results.
- Distribution of aggregable alerts per unique alert is similar to other use cases.
- On the figure bellow we can clearly see that the time difference is often 5 minutes or 1 hour due to default settings if the sensors.

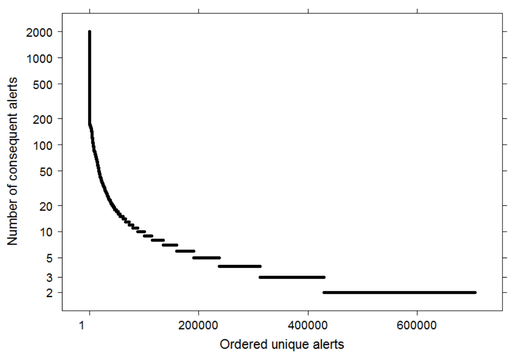

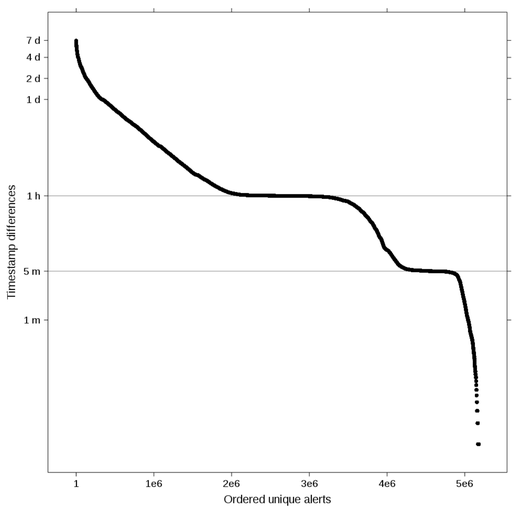

Overall Results

- The table shows numbers and shares of aggregates from all the use cases.
- In total, we identified over 85 % of alerts to be aggregable, but the distribution is not uniform.

	Count	Share [%]
Aggregable alerts	**6,915,568**	**85.98**
Duplicates	99,437	1.24
Consequent	4,379,356	54.45
Overlapping sensors	34,612	0.43
Non-overlapping sensors	2,402,163	29.86
Unique alerts	**1,127,910**	**14.02**
Aggregated	889,853	11.06
Not aggregated	238,057	2.96
Total	**8,043,378**	**100**

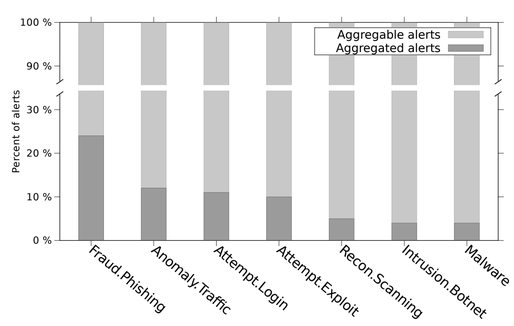

- The portion of aggregable alerts differs between various alert types. The detailed breakdown of alert types is on the following figure:

Experiment and Dataset

- 8,043,378 real-life alerts from SABU alert sharing platform,
- obtained during one week in June 2016,
- 25 sensors, 7 organizations, 1 third-party reporter, 17 alert types.

	Source IP	Source Port	Destination IP	Destination Port	Event Type	Sensor ID	Timestamp
Duplicates	=	=	=	=	=	=	=
Continuing Alerts	=	irrelevant	=	=	=	=	>
Alerts from overlapping sensors	=	=	=	=	=	≠	≃
Alerts from non-overlapping sensors	=	irrelevant	≠	=	=	≠	≃

Use Case 3 – Overlapping sources

- Overlapping sources are sensors that overlap in their detection scope, e.g., sensors A and B from the picture:

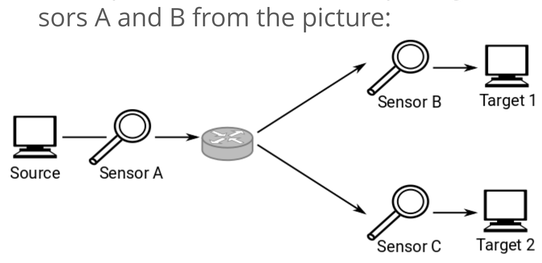

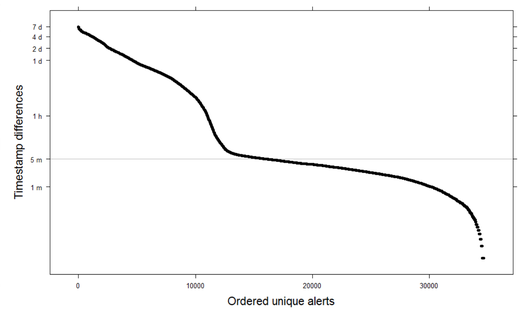

- Typically, the same event is reported by a sensor in a campus network and simultaneously by a sensor in backbone network.
- Similarly to Use Case 2, we can see that most of the time difference are around 5 minutes.

Use Case 4 – Non-overlapping sources

- Similar to Use Case 3, the event is reported by multiple sensors, but these sensors have no overlap.
- Varying time differences, but similar long-tail in aggregates per unique alert.

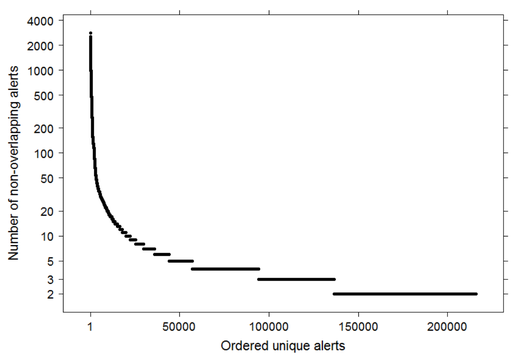

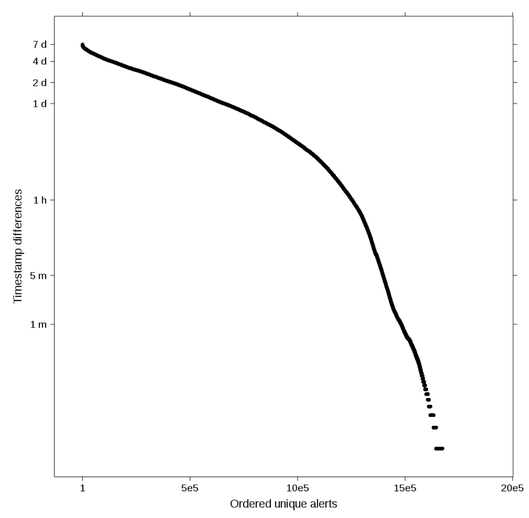

Acknowledgement

This research was supported by the Security Research Programme of the Czech Republic 2015 - 2020 (BV III / 1 VS) granted by the Ministry of the Interior of the Czech Republic under No. VI20162019029 The Sharing and analysis of security events in the Czech Republic.

https://sabu.cesnet.cz/
@CESNET_CERTS

https://csirt.muni.cz/
@csirtmu

CSIRT-MU