INTERACTIVE PROCESS TREE ANALYSIS

visually explore the behaviour of processes

ROBERT-CARL RAPP, CHRISTOPH MÜLLER, FRANZISKA BECKER, PAOLO PALUMBO, THOMAS ERTL

MOTIVATION

The analysts in security operation centres (SOC) are confronted with numerous security alerts [1] and ensure they do not miss the most critical of them. Analysts hold a decisive role in cyber-investigation, and the use of visual analytics for their interface between sensors and operators can use their full cognitive potential for analysis. We enable visual analytics to investigate process hierarchy and process behaviour in an interactive process tree visualisation.

GOALS

We observed security operation centres (SOCs) and outlined the following goals for our approach:

- Visualise the temporal development of processes together with the hierarchy.
- Makes additional process-related events accessible to analysts.
- SOC analysts have less time to familiarize themselves with radically new visualisation types, therefore we keep the solution close to the currently used indented tree view.

ANALYST WORKFLOW

The analysis begins with an alert from the detection pipeline that classified a process as suspicious.

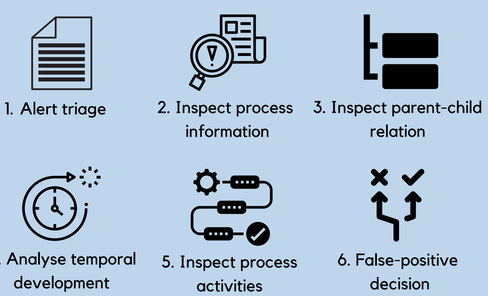

A VISUAL ANALYSIS IN THE PROCESS TREE

1 The visualisation combines a tree layout for hierarchical information on the left and a timeline to explore events and temporal changes. Timeline trees by Burch et al. [2] share a similar layout to view changes over time. The x-axis is overloaded with the hierarchical indentation and time to make interrelated patterns between these dimensions accessible.

The SOC analyst begins the investigation by entering the alert's global unique id to request the data from the backend. The data is collected per-process by endpoint sensors of a commercial managed security product accordingly pre-processed.

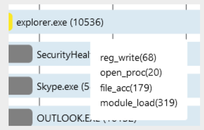

Shows the type and number of process's activities present in the process by hover over the name.

2 Suspicious process names or unusual positions can be the first hint to take a closer look. The analyst can hover over the name to see the summary of activities per process or click on it to view its process details.

All details of a process can be opened in an expandable JSON view by clicking the name.

ANALYSING THE HIERARCHY

3 The hierarchy is represented by means of an indented tree plot formed by the opaque parts of the process bars and emphasised by the blue lines.

- By using coloured pillows, the analyst sees the number of sub-processes and is guided by the colour which processes are part of the requested branch.
- The process is not part of the branch if the pillow is grey like for "Skype.exe (5868)".
- To avoid long loading times for preprocessing the whole process tree, the user can explicitly open such branches on demand by clicking.

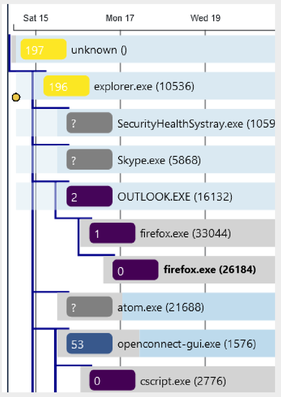

A process branch with five hierarchical levels has coloured pillows and bold text spots out the selected process. In contrast, processes with grey pillows show that they are not part of the suspicious process chain.

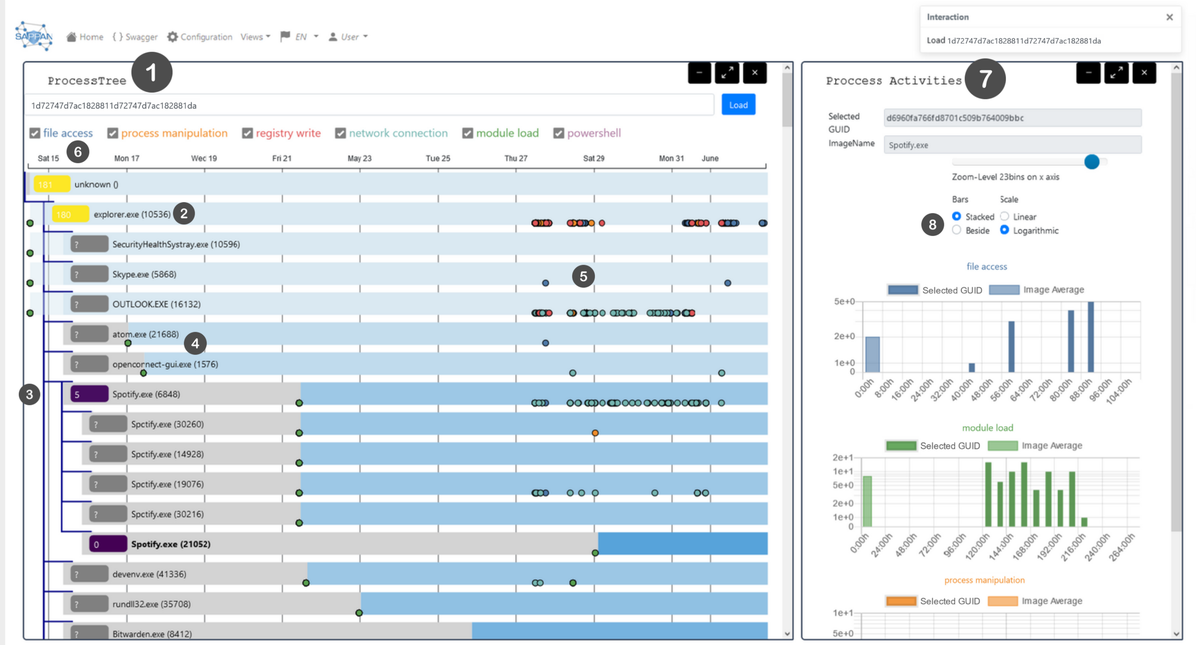

The visualisation is placed in a dashboard composed of cards (ProcessTree, Process Activities) that can be flexibly opened, closed and arranged. Dashboard users can be notified via the upper right message box.

INSPECT TEMPORAL DEVELOPMENT

4 The temporal dimension is mapped on the x-axis and allows the analyst to quickly understand when child processes have been spawned. The process bar is coloured in different shades of blue relative to its process start.

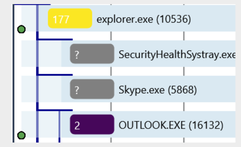

If a process starts earlier than the indentation width for a hierarchy level, its bar is extended left, but using a semi-transparent colour as shown for the three sub-processes.

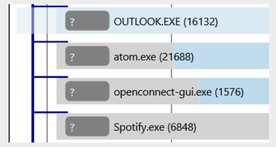

The analysts can use the coloured bars to grasp the start sequence of processes on a hierarchy level. E.g. outlook.exe spawned first followed by atom.exe and openconnect-gui.exe.

Often analysts cannot run their analysis directly after the occurrence of an alert. We included the temporal dimension to allow the analysts to comprehend the endpoint's situation when a suspicious process was detected.

EXPLORE THE PROCESS ACTIVITIES

5 To gain a first overview of what the process did, the process events are represented by dots on its process bar which are colour-coded by the type of the event.

- The legend can be used to filter specific types to avoid overdraw of dots and to see correlations between event types.
- Analysts can hover over a dot to see the full details of the event in a tooltip and raising the dot to the front.

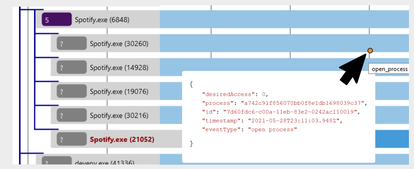

The occurence of open process events highlights the name of the opened process in red like the process "Spotify.exe (21052)" so that analysts can be aware of suspicious call chains.

6

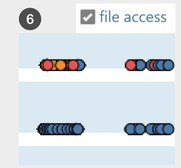

Apply filters to the dots helps the analyst to see specific types like the *file_access* only.

7 While users can see the series of activities related to a process, it is hard to judge whether this time series is reasonable or might indicate a problem. Therefore, it is possible to compare the occurrence in a separate histogram view that shows two histograms per event type:

- One histogram shows the occurrences of events for the selected process.
- The other histogram displays the average of occurrences for all processes of the same process image.
- As time is relative to the start of each individual process, these charts immediately show whether the selected process performed activities at similar points in its life cycle.

8 The view offers linear and logarithmic scaling and interactive zooming to counter the conflict between the histogram covering large ranges and the fact that a single activity might be malicious.

RESULTS

- Support behaviour-based analysis of malicious processes with visual analytics.
- Endpoint sensor data is visually accessible so that analysts can find patterns and anomalies more easily.
- Analysts can view two dimensions at a glance and investigate patterns in the hierarchy and temporal development of process activities.
- The visualisation leans on currently used visualisation in visual interface between sensors and operators makes it possible to deal with novel attack patterns that static methods cannot handle.

FUTURE WORK

- Analytical provenance can capture the interactions and insights of analysts to make the workflow within our application predictable for further reuse.
- Analysts would profit from a filter and search function for process bars to choose a relevant subset of process bars to be displayed.
- Investigating the visualisation utility and usability by a real-world scenario would show various analysts' requirements for our approach.
- Process information enrichment by external security applications like *Cyberchef* would avoid an additional step to the providers' webpage in the analysis workflow.

REFERENCES

[1] F. B. Kokulu, A. Soneji, T. Bao, Y. Shoshitaishvili, Z. Zhao, A. Doupe, and G.-J. Ahn. Matched and mismatched SOCs. In Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, pages 1955–1970, 11062019.

[2] M. Burch, F. Beck, and S. Diehl. Timeline trees. In S. Levialdi, editor, Proceedings of the working conference on Advanced visual interfaces, page 75, 2008.

University of Stuttgart
Germany

ACKNOWLEGMENT

Funded by European Union's Horizon 2020 research and innovation programme under grant agreement No. 833418

CONTACT

Robert-Carl Rapp
robert.rapp@vis.uni-stuttgart.de

Institute for Visualization and Interactive Systems