

BANKSAFE: A Visual Situational Awareness Tool for Large-Scale Computer Networks

VAST 2012 Challenge Award:
Outstanding comprehensive submission, including multiple vizes.

Fabian Fischer*
University of Konstanz
Germany

Johannes Fuchs†
University of Konstanz
Germany

Florian Mansmann‡
University of Konstanz
Germany

Daniel A. Keim§
University of Konstanz
Germany

ABSTRACT

With the reliance of businesses, public institutions and individuals on large computer networks, maintaining their security becomes essential to ensure integrity. To achieve situational awareness, we developed BANKSAFE, which is a scalable, distributed and web-based visualization system to analyze health monitoring data and security datasets. To handle large amounts of data a cloud-based backend database is used to store and analyze the raw data. To evaluate the effectiveness of our approach we use both VAST 2012 mini challenges. Our case studies successfully identify suspicious events, trends and patterns using multiple visualizations.

Index Terms: C.2.0 [Computer-Communication Networks]: General—[Security and protection]; I.3.8 [Computer Graphics]: Application—; H.5.2 [Information Interfaces and Presentation]: User Interfaces—

1 INTRODUCTION

BANKSAFE is a situational awareness application for large-scale computer networks. The overall architecture is shown in Figure 1. To achieve scalability for large datasets BANKSAFE makes use of the cloud-based database service *Google BigQuery* in which monitoring (i.e., health and status checks) and security data, (i.e., IDS alerts and firewall logs) are directly imported. BANKSAFE is a *Java Web Application* hosted by *Apache Tomcat*. To further improve performance and to reduce costs, we make use of high-performance caching systems (*Ehcache*, *Memcached*). Additionally, BANKSAFE provides a web-based graphical user interface using the *Vaadin Java Web Framework*. The visualizations are implemented using *Java Applets*, *HTML5* and *D3.js*.

2 VISUALIZATIONS

Besides standard bar charts to represent the number of active hosts or events, BANKSAFE includes several visualizations to support the analyst in getting an overview, finding trends and identifying suspicious events.

2.1 Visualizations for Health Monitoring

The *Treemap Timestamp Snapshot* visualization (Figure 5) provides a point in time overview of health or policy status for millions of computers. The hierarchy is mapped to the different organizational levels of the company. The sizes of the rectangles are proportional to the number of underlying hosts, while color represents either

*e-mail: Fabian.Fischer@uni-konstanz.de

†e-mail: fuchs@dbvis.inf.uni-konstanz.de

‡e-mail: Florian.Mansmann@uni-konstanz.de

§e-mail: Daniel.Keim@uni-konstanz.de

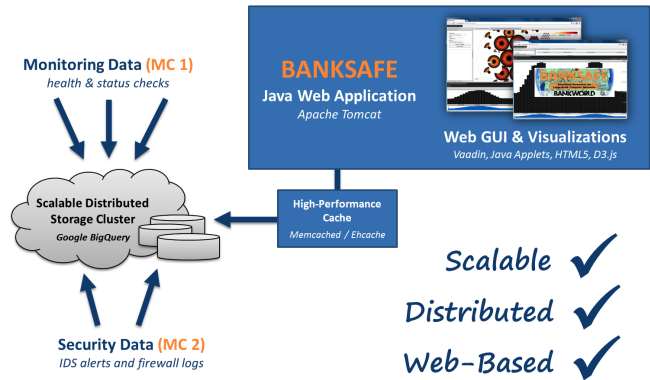


Figure 1: Overview of the BANKSAFE architecture.

health or policy status to reveal facilities with a high percentage of suspicious machines. To provide a compact, yet high-density temporal information display, we additionally implemented a *pixel-based matrix*. This 5x5 colored matrix depicted in Figure 2 represents the number of hosts in color for all possible combinations of policy and activity scores of a single region for an aggregated time-span. By arranging these matrices in a small multiple display, clear temporal patterns can be seen in Figure 4.

2.2 Visualizations for Firewall Logs and IDS Alerts

To visualize time-series data of the firewall log within their respective hierarchy the *ClockMap* [1] visualizations can be used. Basically a circular treemap is used and enhanced with circular temporal glyphs using a clock metaphor to represent 24-hour time-series data as shown in Figure 3. To analyze IDS events, the analyst can use the *Relaxed IDS Timeline* [2]. Each row in Figure 7 represents the events for a particular source IP address and each fixed-size column shows the events of one hour. Color is mapped to the event classification attribute, which helps to visually distinguish the event types. Selecting an event gives more information and highlights all other events of this particular type using connecting lines.




Figure 2: Activity-policy matrix for all hosts in a region within one hour.




Figure 3: Circular clock glyph representing 24-hour time-series for one host.



Figure 4: Small multiple visualization using activity-policy matrix representations.

3 CASE STUDIES

In the following, we briefly discuss some findings¹ for both VAST 2012 mini challenges using BANKSAFE. In the first challenge, the analyst is interested in the network status of the entire *Bank of Money* at a specific point in time. To answer such questions the *Treemap Timestamp Snapshot* can be used as depicted in Figure 5.

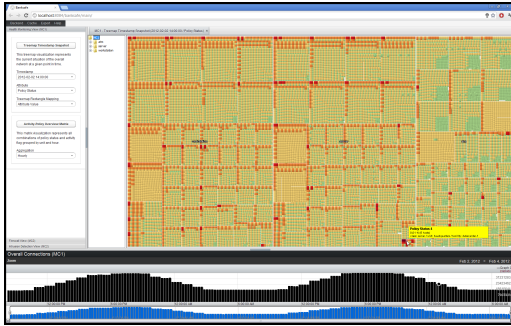


Figure 5: Treemap visualization for detailed point in time overviews.

To spot outliers the analyst can interactively explore the treemap and visually compare different regions. In *region-5* and *region-10*, for example, all machines suffer from moderate policy deviations (no green rectangles at all). Another important analysis question is temporal trend detection, which can be addressed using the *Activity-Policy Overview Matrix*. Figure 4 shows the small multiple visualization representing the overview matrices for *region-5* for the first 23 hours. This pixel visualization highlights that many machines get active during hour 5 to 15. Moreover, there is a continuous shift to a higher policy level in each hour, which means that gradually more and more hosts within this regions become infected.

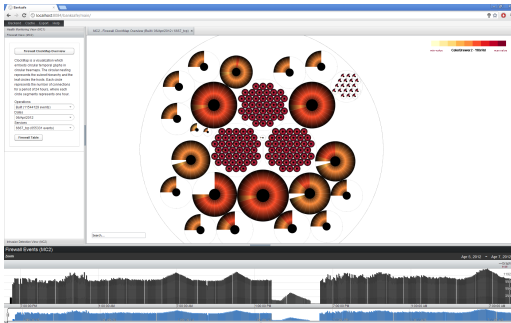


Figure 6: *ClockMap* visualization to identify suspicious hosts.

In the second mini challenge, the analysts use security datasets to identify trends and suspicious events. To visualize time-series data of the firewall log within their respective hierarchy the *ClockMap* visualizations can be used by the analyst. Figure 6 shows the traffic of all source machines, connecting to an IRC service on port 6667/TCP. To explore which subnets start to communicate in which hour with the external IRC servers, we use the 24-hour glyph representation, where each segment of a circle represents one hour.

¹For more details about the findings, please refer to the supplementary material (video and challenge submission).




Figure 7: *Relaxed IDS Timeline* to explore IDS alerts in time.

The *Relaxed IDS Timeline* in Figure 7 shows that several hosts are producing IRC authorization messages. It seems that those machines became suddenly infected and attempted to talk with their bot master over IRC. With the help of BANKSAFE, the administrators can query for the destination IP addresses of the remote IRC servers. Blocking those IP addresses on the firewall would prevent the infected machines to load additional malware or to react to new commands by the bot master.

4 CONCLUSIONS

To enhance security and to provide situational awareness for large-scale computer networks BANKSAFE uses a combination of multiple visual representations. Figure 8 shows exemplarily how BANKSAFE can be applied to control room situations to analyze and present analysis results in a scalable way. Through these multiple cross-linked visualizations several trends, patterns and suspicious events have been successfully identified in both datasets of the VAST Challenge 2012.



Figure 8: BANKSAFE used in a control room.

ACKNOWLEDGEMENTS

Parts of the research leading to these results has received funding from the European Commission's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 257495, "Visual Analytic Representation of Large Datasets for Enhancing Network Security" (VIS-SENSE).

REFERENCES

- [1] F. Fischer, J. Fuchs, and F. Mansmann. ClockMap: Enhancing Circular Treemaps with Temporal Glyphs for Time-Series Data. In *Proceedings of the Eurographics Conference on Visualization (EuroVis 2012)*, 2012.
- [2] F. Fischer, F. Mansmann, and D. A. Keim. Real-Time Visual Analytics for Event Data Streams. In *Proceedings of the 2012 ACM Symposium on Applied Computing (SAC 2012)*, 2012.