Integrated Tool Support in the Context of Security and Network Management Convergence

Felix von Eye, David Schmitz, and Wolfgang Hommel
Leibniz Supercomputing Centre, Munich Network Management Team
Garching n. Munich, Germany
Email:{voneye,schmitz,hommel}@lrz.de

***Abstract*—The convergence of different types of networks, such as for telecommunication and data transfer, on the hardware layer has an obvious impact on both network management and security management, which especially affects network service providers as well as data centers. This paper argues that also methods, algorithms, and tools from both research domains, network management and security management, should systematically be reviewed for synergies. It first analyzes the current state of the art in both domains and identifies gap areas that require further investigation. Then, the Customer Network Management for the X-WiN (WebCNM) network management tool, which started as a research prototype in the pan-European research and education network, GÉANT, is presented along with selected extensions that were designed and implemented to integrate security management functionality. Several security event visualization options and their use within the European industry-focused Safe And Secure European Routing (SASER) project are discussed.**

Keywords–Network management; Security management; Integrated management; Enterprise management; Convergence.

I. INTRODUCTION

In both research and real-world operations, network management and security management are often treated as complementary but still largely independent parts of the overall IT service management performed by network service providers and data centers. However, this view limits the prospects in case of, e.g., handling incidents. In many situations both security events and network events are related and visible within the same time interval, for example when a security event results in an unusual or peak usage of the network. In turn, certain conditions, such as an unreachable service or device, can be caused either by a denial of service attack or a fault in the network.

As security management and network management are typically supported by independent technical management software tools, such contextual relationships are not visible without further ado. This motivates a new approach to combine the information processed by either management tool landscape into one, leading to a convergence of selected network and security management methods and tools. Although it is obvious that it is neither reasonable nor practically possible to fully integrate both disciplines currently, we elaborate on selected event categories that result in operational benefits or have interesting properties for research that motivate further investigation.

The work leading to the presented results have been carried out in the large-scale distributed environment of the SASER-SIEGFRIED project (Safe and Secure European Routing) [1], in which more than 50 European research and industry project partners design and implement network architectures and technologies for secure future networks. The project's overall goal is to remedy security vulnerabilities of today's IP layer networks and have them ready for deployment in backbone networks by the year 2020. Thereby, security mechanisms are designed based on an analysis of the currently predominant security problems in the IP layer as well as upcoming issues, such as vendor backdoor and traffic anomaly detection. The project focuses on inter-domain network traffic and routing decisions that are based on security metrics, which are derived from aggregating and combining security measurements carried out by multiple involved organizations in a cooperative manner. As this scenario relies on the future internet architecture and software-defined networks (SDNs), we present a prototype of a combined security and network management system, which is independent of any commercial vendor or operator and has already been used in the context of GÉANT, the pan-European research network.

The remainder of this paper is structured as follows: Section II gives an overview of the current state of the art and related work regarding network management, security management, and the convergence of these two disciplines. Section III presents the WebCNM framework, which serves as technical basis for implementing an integrated network and security management platform, and experiences made in the SASER project. Finally, Section IV concludes the paper and gives an outlook to future work.

II. STATE OF THE ART

Both network management and security management are computer science disciplines with a very long, yet only partially overlapping history. In this section, first the focus is clarified by discussing the term *management*. Then, we analyze the status quo of network management in Section II-A and the current state of the art of security management in Section II-B. Finally, Section II-C reviews related work on the convergence of both disciplines and outlines how it has influenced the improved approach presented in Section III.

In general, *management* in the context of IT services refers to any measures and activities that are performed in order to achieve effective and efficient operations of those IT services and the required resources in alignment with an organization's business goals [2]. Management therefore covers the whole life-cycle, including planning, provisioning, setup, configuration, operations and maintenance, and removal; it involves personnel, procedures, processes, technology, and software tools. Management must be performed on any abstraction layer, such as individual physical hardware components or software

 ISBN: 978-1-61208-416-9

properties (e.g., application response time) up to an enterprise-wide and even inter-organizational view. The term *integrated management* refers to approaches that successfully deal with heterogeneity, such as managing hardware by different vendors or across various types of systems, such as network components, servers, and application software. Usually, *management architectures* describe various properties of how management is carried out in an abstract manner. They can be broken down into four models:

1) The systems to be managed, referred to as *management objects*, are described by an *information model*.
2) The roles of all systems involved in management along with their types of cooperation are described by an *organizational model*.
3) The *communication model* describes the exchange of management-related messages between the roles defined by the organizational model.
4) The *function model* groups and structures the management-specific functionality.

A software implementation of a *management architecture* is referred to as a *management platform*. Furthermore, a *management system* refers to the sum of organizational structure, policies, planning activities, responsibilities, practices, procedures, processes, and resources [3]; i.e., a management system is not just a piece of software, but management software, such as a *management platform* and various *management tools*, contributes to this overall management system.

A. Network Management Status quo

Network management has evolved over decades with the growth of intra-organizational enterprise networks and the Internet. The most influential foundation for network management is the classical OSI management architecture, which proposed the five functional areas fault, configuration, accounting, performance, and security (FCAPS) [4].

Fault management, which is closely related to the IT service management areas of incident management and problem management as used by the standard ISO/IEC 20000, deals with anticipating, detecting, and reacting to any types of network faults, such as hardware defects, resource exhaustion, and quality of service guarantee violation. To a large degree, fault management is based on monitoring one's network using both active and passive measurements. Inter-domain fault management is still non-trivial due to the heterogeneity of the involved hardware and organizations' restrictive information sharing policies [5].

Configuration management is the function area to actively modify a network component's parameters. It is an area for which integrated management has not been established successfully because of the inherent complexity of the task:

- Various different types of network components, such as routers, switches, and WiFi access points must be supported; they operate on different layers of the ISO/OSI model.
- Hardware and firmware heterogeneity: The same type of network component, such as a router, often has vendor-specific configuration options and still shows a lack of compatibility when components from different vendors or product generations are mixed.
- Scalability: Larger organizations typically must manage several thousands of network components; however, many management tools still lack support for parallelized operations, e.g., to perform even simple tasks like firmware updates.
- Emerging technologies: Given the cost of network equipment, it is not unusual that network components are in use much longer than other types of hardware. For example, while modern network equipment supports new management paradigms, such as SDNs, the majority of components in use today still needs to be managed using Simple Network Management Protocol (SNMP).

As a consequence, network component configuration management is mostly done by using vendor-specific tools in practice, which means that organizations need other means to ensure configuration consistency across network component types and hardware models.

Accounting, like fault management, is based on monitoring network components. Various low-level hardware counters are aggregated and combined to support, e.g., billing processes.

Performance management is closely related to the management of service levels. Quality of service parameters typically include the minimum guaranteed availability of network links along with their bandwidth and upper limits for undesired properties, such as packet loss, delay, and jitter. Similarly to accounting, these parameters may vary with the actual content that is being transported, because, for example, voice-over-IP connections have different requirements than bulk data transfers.

Security management, as far as the traditional network management is concerned, focuses on security properties of individual network components, such as authentication and authorization for management access. While most modern network components support some additional security features, such as port-based access control in switches and access control lists in routers, these features either depend on additional central security management components, such as a RADIUS server for user or device authentication, or are clearly limited by their hardware performance.

B. Security Management Status quo

The classic goal of security management is to ensure the confidentiality, integrity, and availability of sensitive data as well as the systems and services that process it. As in other areas of computer science, there is no single silver bullet security measure to achieve this goal, but numerous modular security controls must be combined in a complementary and partially deliberately overlapping manner, which is often referred to defense-in-depth and graceful-degradation. Security controls can be categorized as *preventing*, *detecting*, or *reacting* in relation to successful attacks.

Derived from the work by Hyland and Sandhu [6], the following areas of security operations must be considered:

- *System security* covers the management of available software updates, hardening each individual system, having malware protection in place, integration with central security services, such as authentication servers, and preventing undesired data leakage.

 ISBN: 978-1-61208-416-9

Network components are also considered to be such systems.

- *Network security* encompasses the definition of security zones based on protection requirements, the application of virtualization and segregation concepts, such as virtual local area networks (VLAN), and dedicated network security components, such as firewalls and network-based intrusion detection systems.
- *Access management* covers the management of users along with their roles and permissions as well as setting up authentication and authorization services. Privileged accounts, i.e., those used by administrators with full control over a system, deserve special attention. Physical access to systems also must be considered.
- *(High) availability* is practically only achieved through some type of redundancy, which includes both hardware and copies of the data.
- *Cryptography* is a complex discipline with endless applications in information security. Flawed implementations and improper application due to lacking know-how are also a large source of major security problems.

Preventive security controls, such as firewalls and access control mechanisms, are intended to enforce policies, i.e., someone must define what is allowed or undesired in an a-priori manner. The proper implementation of these controls can be checked, e.g., by means of penetration tests. The safe general assumption is that at least parts of one's IT infrastructure are insecure and the methods provided by the subdiscipline of vulnerability management assists in identifying those parts. To prioritize options for the improvement of the overall security level, risk management methods need to be applied.

However, security monitoring is quite mature. It mostly relies on aggregating, correlating, and evaluating security events provided by various dedicated sensors as well as systems and applications, e.g., via log files. Intrusion detection systems passively monitor data as it is processed by the involved systems and use signatures of known attacks or outlier detection to register unusual behavior. As most attack attempts are background noise, i.e., they are not specifically targeted against systems known to be vulnerable, correlation with an asset management database and information from one's own vulnerability management greatly help to reduce the number of false positive alerts. Security information and event management (SIEM) systems perform this correlation, prioritize detected security incidents, create reports, and can be considered to be the security management counterpart of management platforms in network management.

C. Related Work on Management Convergence

It is obvious that network management and security management have partially overlapping scopes: Certain network components, such as firewalls, are typically operated by IT security personnel, and security events generated by network components are being processed by security management tools as a matter of course. However, this information flow can traditionally be considered one-way, from network management to security management tools, without feedback loops and with completely separated tool sets used.

Convergence of both network and security management has the meaning that methods, procedures, and tools can jointly be used for both disciplines, bringing them together as one to increase effectiveness and reduce the overhead caused by separate processing of the same data in multiple instances. In this section, related work in this area is analyzed, which has influenced the design of the approach presented in the next section.

In [7], Dawkins et al. presented a novel network security management system for tracking large-scale, multi-step attacks based on data from various specific sensors in different networking domains. Their system provides real-time correlation and analysis of the data; they focus on an interesting novel visualization method that provides a kind of heads-up cockpit display of an entire network that can be used by network as well as security management personnel. While their work is intended for converged networks and sensors specifically designed for such environments, our approach differs in that it makes use of existing data sources.

In [8], Kuklinski and Chemouil discuss management challenges specifically for SDNs. The propose a mapping of the classic FCAPS approach and point out the special role of SDN controllers, for which additional security monitoring mechanisms are proposed. Complementary, Zhu et al. discuss the role of vendor specifications in cross domain communication in [9]. They present key requirements for an inter-domain security infrastructure along with a reference architecture.

Han and Lei compare the policy languages that are used for network and security management in [10]. The identified similarities make it interesting to formulate policies and rule sets that cover both management disciplines. However, despite some widely used policy languages, the use of proprietary policy specification formats in various products still fuels the demand for inter-system conversions. Wang-fei and Qi developed a novel network management system in [11] with an emphasis on security management that covers networking equipment as well as virtual machine servers; it also addresses combined performance management and unifies the access control mechanisms of both network components and IT services.

In sum, despite several approaches towards combining both management disciplines in prior work, there is no thorough analysis of which types of network and security events would be relevant for a unified management approach and no integrated management platform exists yet. This motivates the approach discussed in the next section.

III. Extending the WebCNM Framework

WebCNM started as a research prototype for vendor-independent, multi-domain, and customer-oriented network monitoring and management visualization. Its core functionality consists of the visualization of network maps that are organized in a tree-based hierarchy. Each network map shows network elements, i.e., nodes and links of different network layers, together with current or historic status and metric information. Detailed historic statistics are provided in a drill-down manner.

WebCNM was originally developed in the pan-European research and education network, GÉANT, as part of the GN3 project (2011-2013). It has been used to visualize

 ISBN: 978-1-61208-416-9

network topologies and respective network metrics of many European and Non-European national research and education networks (NRENs) including DFN (Germany), SWITCH (Switzerland), Uninett (Norway), Pionier (Poland), GARR (Italy), Surfnet (Netherlands), Renater (France), Hungarnet (Hungary), MRen (Montenegro), GRNet (Greece), SEREEN (South-East-Europe), ESNet (Energe Sciene Network, US), Internet2 (US), RedClara (Pan South/Middle America), RNP (Brazil), and the pan-European core-network of GÉANT as well as the optical private network of CERN's Large Hadron Collider LHC project (LHCOPN). A significantly extended version is in production as of 2015 for all higher education institutions connected to the German NREN. WebCNM's functionalities comprise, among others, customers' network service information, network access information/status, network access accounting, network core status, and network performance management.

WebCNM was specifically designed to be extensible in an flexible and modular manner: It features a JavaScript extension API, which is independent of the implementation technology of both the backend and the GWT-based web client, which allows for a client-side integration with other web pages and web tools. It is therefore a suitable basis for integrating functionality related to security management.

A. Integration of Security Events in WebCNM

Before the integration of security events into any network management tool, it is necessary to analyze which potential security events have an impact on network connections or devices. As a first step, only direct attacks on network devices, attacks on the network protocols or on the availability of systems, services, and the network itself are considered.

As there are a lot of possible attack scenarios, the following implemented examples highlight some core issues:

- *Port scans*: A port scan is often treated as an attack in which a client attempts connections to a range of server ports with the goal of finding active services and preparing the exploitation of their known vulnerabilities.
- *Border Gateway Protocol (BGP) Hijacking*: BGP hijacking is the illegitimate takeover of groups of IP addresses by corrupting inter-domain routing tables. This attack has a very high impact on network security but it is difficult to implement as the attacker has to first compromise central network infrastructure components, such as the backbone routers.
- *Amplification Attack*: In distributed reflective denial-of-service (DRDoS) attacks, adversaries send requests to public servers (e.g., open recursive DNS resolvers) and spoof the IP address of a victim. These servers, in turn, flood the victim with valid responses and – unknowingly – exhaust its bandwidth.
- *Backdoor*: A backdoor in a computer system is a method of bypassing security controls, such as user authentication, in order to enable unsolicited remote access to a system, obtaining access to data, and so on, while attempting to remain undetected. While backdoors on servers and workstations are quite usual and a well researched topic, backdoors on switches and routers are still a huge problem.

Out of these threat scenarios, denial-of-service attacks are the most simple example of security incidents that are also visible in network management because the bandwidth utilization parameter increases. To enable the correlation between network and security events, WebCNM was extended with an interface to use intrusion detection message exchange format (IDMEF) based messages [12], which are designed to exchange security relevant data between systems and domains.

B. Visualization of logins on network devices

In this paper, an example is used to clarify the ideas behind the new approach, in which logins to network devices like backbone routers are analyzed. Regular secure shell (SSH) logs are used as basis, which gives the possibility to detect whether a suspicious login has occurred.

In WebCNM, any network device can be monitored by the network management tool part. There are two ways to determine if there is a security incident. The first one is the alerting by the SIEM system via the IDMEF interface. In this case, the SIEM system sends an IDMEF message as a notification to WebCNM and this event is then displayed inside the network management tool. This method is usually used to inform the network administrators about suspicious events in the network respectively on network devices. In general, in a regular sized network there are too many events, so they can be only analyzed by visualization. In this case, the affected subnetwork or network device is highlighted inside the network management system, which helps to find possible network issues.

Inside the SIEM system, the security administrators usually use a lot of different rulesets, which are able to automatically detect attacks or misuse of components inside the network. These rulesets have to be adjusted to produce only a few false positives and negatives. To detect unknown attacks or to detect events, which are slightly under the radar, it is necessary to also manually analyze the communication flows.

This leads to the other way to display security-relevant events during the direct inspection of the network components. For example, the log messages of a network device can be displayed in the network management tool. As the network management tool is not a SIEM system, it is not able to process the correlation of events in an automated manner, but instead can only visualize them for the administrator. This allows network administrators to see suspicious events, e.g., a login to the network device from outside the management VLAN, which is quite unusual. Figure 1 illustrates an example of a possible visualization of SSH logins that has been developed within the SASER project: It uses GeoIP-based grouping of remote source addresses and makes it easy to visually distinguish between successful and failed logins. Brute force attempts to guess user passwords can be spotted on the right, and clicking on any edge or vertex brings up more details about the selected group of events. When viewing more generic events, advanced color schemes can be used, e.g., to visually identify protocols or VLANs. In this scenario, the server is only accessible via the SSH public key method for the administrators, so it is unusual that there are also connections authenticated with passwords. In other cases this behavior is completely different, as the authentication via password is the only possible way to access a device. This diversity of authentication methods prevents a fully automated generation of SIEM rules, e.g., for anomaly

 ISBN: 978-1-61208-416-9

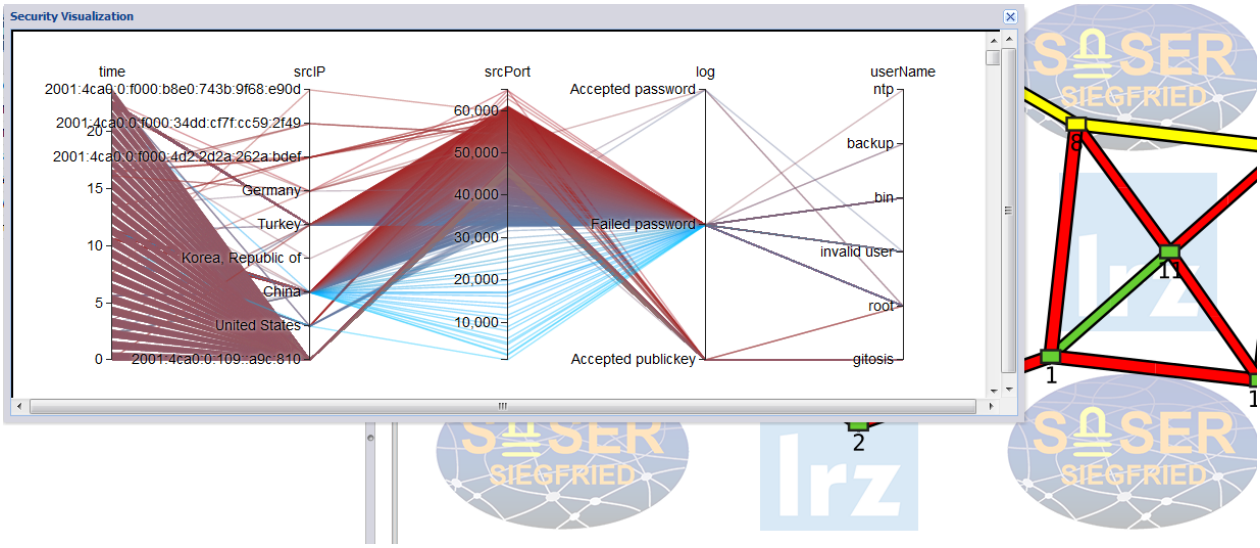

Figure 1. Visualization of SSH logins in WebCNM.

detection, but instead needs manual classification. This can be done very user-friendly via visualization.

As the information collected in such a way is only some type of metrics from the WebCNM point of view, the used visualization is not limited to the form shown in Figure 1. Figure 2 shows some visualizations that were developed together with the user interface division of FH Potsdam in the SASER project. Especially for data for which no automated processing and alerting rules have been implemented, these visualizations exploit the innate human skill to quickly identify patterns and outliers that do not match one of these patterns; therefore, while many of those developed visualizations do not provide explicitly the necessary information to handle the potential incidents they indicate, they nevertheless assist in quickly picking up those events manually that have to be reviewed further.

These different visualization types were developed, as there are no visualization types generally applicable to all attack scenarios or systems. The benefit of the human analysis of the visualizations is that humans can browse through different visualization views. If there are any outliers visible, they are able to correlate them with other outliers in other views enhanced with their knowledge. Therefore, it is important that there are ways to switch the kind of visualization. The visualizations shown in Figure 2 are specially made for detection of amplification attacks, denial of service detection, portscan detection, the correlation between source and destination addresses in login events, the delays of Transmission Control Protocol (TCP) connections, and the detection of high traffic network nodes. Often, the design goal of visualizations is to highlight differences or similarities of the received datasets. This helps the network administrators to detect security events before the attacker becomes visible to the security administrators and their SIEM system.

C. Evaluation in the SASER Scenario

In the SASER project, the newly designed network management system was introduced to allow the integrated view on network and security. The implementation was tested by several partners to identify whether the concept also works in real-world scenarios. It turned out in these tests, however, that the focus on port scans is not very useful for huge Internet service providers as there are too many messages to process, which makes the overview very difficult to keep. But they proposed the small change that the messages of lesser important events should only be visible if an administrator is searching for a fault.

On the other side the possible amplification vulnerabilities turned out as very useful. For visualization, there are in general too many events, but it turned out that especially for connections with very strict service level agreements it is useful to know if there is a potential denial of service threat by a vulnerable device inside the local domain, so the routing can be set to minimize the risks.

The function to analyze the log messages of a router or switch with regard to security was determined to be very helpful, as this work is done by security staff, which often has limited knowledge about network-side topology changes, or requires to keep track of those changes manually. Both alternatives are based on the fact that the root causes of faults are hard to determine.

Furthermore, the project's focus on SDN enables a lot of new possibilities in monitoring and management. As the SDN controller is designed to get extended with applications, it is possible to connect the monitoring functions directly to the controller. This leads to more efficient algorithm implementations because the export and transformation of the information is not needed anymore. There are also no additional delays between the monitoring and the analyzing. As the network management system is directly connected to the SDN controller, it is also possible to get the information directly from the controller.

IV. Conclusion and Outlook

As the integration of security management functionality, including the visualization of security events, as an extension to the WebCNM research prototype has shown, there is an in-

 ISBN: 978-1-61208-416-9

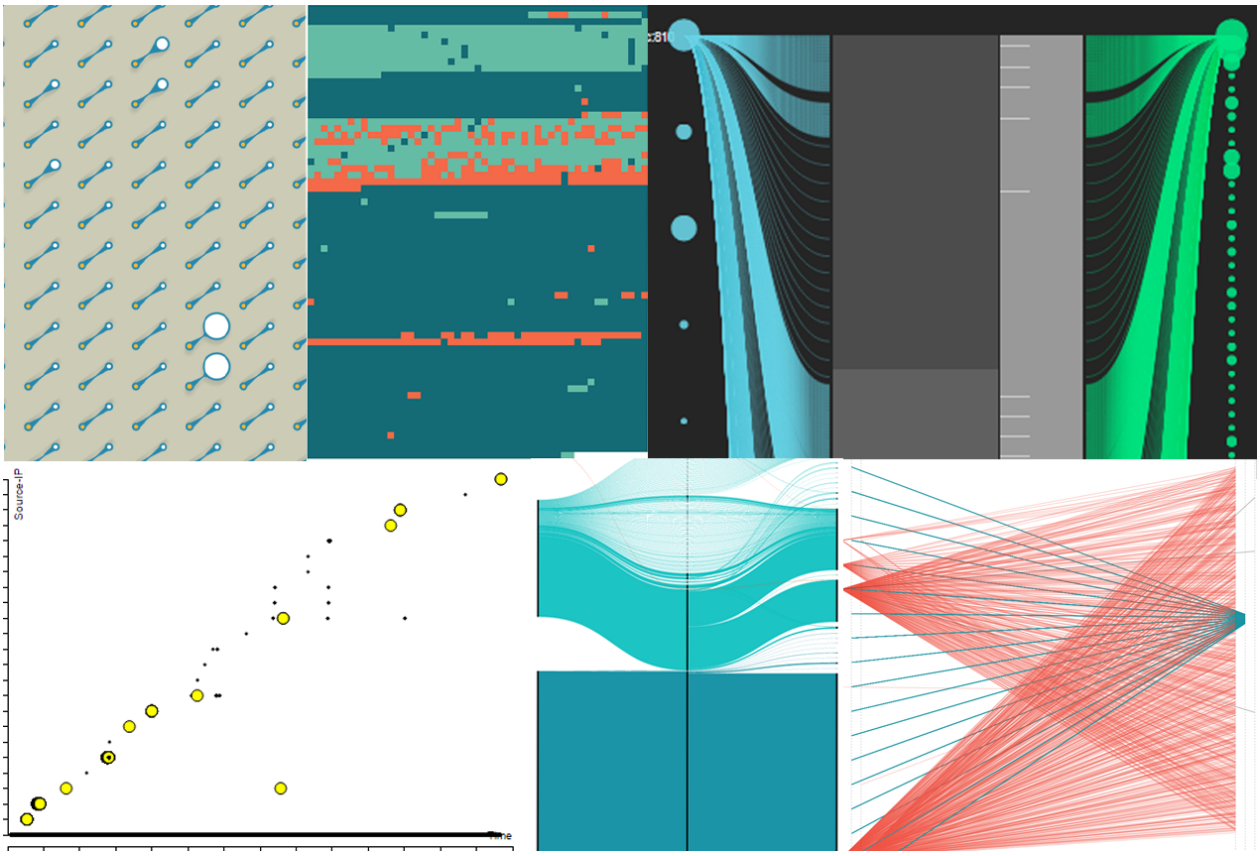

Figure 2. Different visualizations of security events in WebCNM.

teresting potential for converged network and security management tools. Although decisions about which events to include and how to visualize them still require more research, it is obvious that future management software suites can cover both disciplines in integrated manner. While presently only security events from data sources within an organization have been processed, our future work will investigate the inclusion of events and configuration items created by SDN controllers and so-called SDN applications with the goal of enabling WebCNM-based network and security management across organizational borders. The primary security management challenges to this extent are restrictive information sharing policies and technical heterogeneity in the real world, similar to the previous network management challenges that have successfully been overcome. In the long term, self-adapting monitoring data sources that automatically adjust, e.g., their threshold parameters before raising alarms, will also be integrated into WebCNM.

ACKNOWLEDGEMENTS

Parts of this work has been funded by the German Ministry of Education and Research (FKZ: 16BP12309). The authors wish to thank the members of the Munich Network Management (MNM) Team for helpful comments on previous versions of this paper. The MNM-Team, directed by Prof. Dr. Dieter Kranzlmüller and Prof. Dr. Heinz-Gerd Hegering, is a group of researchers at Ludwig Maximilian University of Munich, Technische Universität München, the University of the Federal Armed Forces, and the Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities.

REFERENCES

[1] "The SASER-SIEGFRIED Project Website," http://www.celtic-initiative.org/Projects/Celtic-Plus-Projects/2011/SASER/SASER-b-Siegfried/saser-b-default.asp, 2013, [retrieved: 2015-04-10].

[2] H.-G. Hegering, S. Abeck, and B. Neumair, Integrated Management of Networked Systems — Concepts, Architectures and their Operational Application. ISBN 1-55860-571-1, Morgan Kaufmann Pub., 1999.

[3] ISO/IEC 27001:2005, Information technology – Security techniques – Information security management systems – Requirements. International Organization for Standardization and International Electrotechnical Commission, 2005.

[4] ISO/IEC 7498-4:1989, Information processing systems – Open Systems Interconnection – Basic Reference Model – Part 4: Management framework. International Organization for Standardization and International Electrotechnical Commission, 1989.

[5] P. Marcu, D. Schmitz, W. Fritz, M. Yampolskiy, and W. Hommel, "Integrated monitoring of multi-domain backbone connections," International journal of Computer Networks & Communications (IJCNC), vol. 3, no. 1, 2011, pp. 82–99.

[6] P. C. Hyland and R. Sandhu, "Management of Network Security Applications," in Proceedings of the 21st NIST-NCSC National Information Systems Security Conference. Arlington, Virginia, USA: National Institute of Standards and Technology, Oct. 1998, pp. D–86–D–97.

[7] J. Dawkins, K. Clark, G. Manes, and M. Papa, "A Framework for Unified Network Security Management: Identifying and Tacking Security Threats on Converged Networks," Journal of Network and Systems Management, vol. 13, no. 3, 2005, pp. 253–267.

[8] S. Kuklinski and P. Chemouil, "Network Management Challenges in Software-Defined Networks," IEICE Transactions on Communications, vol. E97-B, no. 1, 2014, pp. 2–9.

[9] W. Zhu, L. Vizenor, and A. Srinivasan, "Towards a Reference Architecture for Service-Oriented Cross Domain Security Infrastructures,"

 ISBN: 978-1-61208-416-9

in Internet and Distributed Computing Systems, ser. Lecture Notes in Computer Science, G. Fortino, G. Di Fatta, W. Li, S. Ochoa, A. Cuzzocrea, and M. Pathan, Eds. Springer International Publishing, 2014, vol. 8729, pp. 275–284.

[10] W. Han and C. Lei, "A survey on policy languages in network and security management," Computer Networks, vol. 56, no. 1, 2012, pp. 477–489.

[11] W. fei Zhao and Q. Wang, "Study on Network Management and Security Access Control," Communications Technology, vol. 3, 2011, pp. 93–95.

[12] H. Debar, D. Curry, and B. Feinstein, "The Intrusion Detection Message Exchange Format (IDMEF)," The Internet Engineering Task Force (IETF), Request for Comments 4765, Mar. 2007, [retrieved: 2015-04-10]. [Online]. Available: http://tools.ietf.org/html/rfc4765

 ISBN: 978-1-61208-416-9