Research Statement

Birhanu Eshete

My research interests span the areas of systems security, cyber-crime analysis, big-data security analytics, and machine learning for security. In systems security, I particularly focus on the analysis and detection of advanced and persistent threats, web application security, and web-borne malware defense. In cyber-crime analysis, I focus on malicious sites/URLs, exploit kits, and ransomware.
My research so far has contributed techniques, tools, and experimental insights in: web-borne malware defense [4, 7, 8, 10, 12], counter-offensive cyber-crime analysis [5], automated exploit generation [1], web application misconfiguration auditing [9, 11], and analysis and detection of stealthy cyber-attacks [13]. My work leverages theoretical and practical insights from a range of techniques such as machine learning [7, 8, 10, 12], static analysis [5, 10, 11, 12], dynamic analysis [1], emulation and sandboxing [7, 10], graph theory [1, 8, 13], SAT solvers [1, 5], genetic algorithms [12], information flow analysis [1, 13], and big-data analytics [8, 13].
A common thread in my research is a systematic inference of insights from artifacts produced by systems under nominal and adversarial settings. I then analyze, design, implement and evaluate defense/offense solutions to thwart cyber threats. My core focus in all my research is to come up with solutions that cope with the ever-evolving threat landscape without compromising precision and responsiveness of a cyber threat detection and analysis system. Next, I will give a summary of my research philosophy, past research, and future research directions.

Research Philosophy

The "why" matters: When doing research, there is an often ignored truth that research should be fun to do and it should not be a rote technical gymnastic to comply with academic routine. I believe as much as the "what" and the "how" of doing research matter, so does the"why".
Sharing is crucial: Over the years, I have immensely benefited from research source code and datasets made publicly available by many researchers. I am, therefore, a strong proponent of sharing research insights, source code, and datasets with the research community. Sharing saves one from re-inventing the wheel and leaves enough time to think through the problem at hand.
Aiming for practical application: I believe in building systems that solve important problems. But that is just the beginning unless built systems are field-tested for their practical effectiveness. I believe the best way to build usable and robust solutions is using operational feedbacks from deployed solutions and refine them in iterations.

Past Research

My past research covered web-borne malware, cyber-crime analysis, advanced and persistent threats, web security, and pervasive computing with an emphasis on anomaly detection and context-awareness.

Web-borne Malware: Analysis, Characterization, and Detection

The focus of my Ph.D. dissertation [3] was automated analysis, characterization, and detection of malicious activities on the Web. My work addressed three problems; (i) specific analysis techniques that fail to capture the holistic view; (ii) detection models that couldn't cope with threat evolution, and (iii) the emergence and prevalence of exploit kits.
Holistic and Lightweight Analysis and Detection: To address (i), I developed a holistic and at the same time lightweight approach that effectively detects malicious web pages and demonstrated the technique through a system called BINSPECT [10]. The approach addressed an important

limitation of prior malicious web page detection techniques, i.e., the focus on specific attack payloads. BINSPECT effectively and efficiently combines *static analysis*, *minimalistic emulation*, and *machine learning* to holistically characterize and detect malicious web pages that infect (with malware) systems of unsuspecting visitors. In a large-scale evaluation, BINSPECT achieved 97.8% accuracy with negligible false positives.

Evolution-Aware Analysis and Detection: To address the evolution of attack payloads in both benign and malicious web pages, and accordingly tune detection models, I developed a system called EINSPECT [12]. The core insight in EINSPECT is *evolutionary searching and optimization* via *genetic algorithms* to automatically select the best permutation of features and learning algorithms as web page artifacts evolve over time. Experimental evaluation demonstrated that EINSPECT reduces false positives by about 10% compared to ad-hoc retraining of models.

Behavior-Centric Fingerprinting and Detection of Exploit Kits: Soon after I developed BINSPECT and EINSPECT, exploit kits emerged as the major players in web-borne malware infection. To respond to the threat posed by exploit kits, I developed WEBWINNOW [7], the first behavioral fingerprinting and detection system for malicious URLs that point to exploit kits. WEBWINNOW, leverages *attack-centric* and *self-defense* characteristics of exploit kits to effectively fingerprint and ultimately detect them in real-time. The approach leverages *honeyclients* to probe suspicious URLs in real-time and "milk" behavioral artifacts indicative of exploit kit behavior. WEBWINNOW achieved 99% detection accuracy with minimal false positives, demonstrating its practical effectiveness.

From Offline Infection Analytics to On-the-Wire Malware Detection: Recently, I approached web-borne malware detection from a different perspective. By carefully studying successful malware infection episodes, I modeled web sessions as conversations between a potential victim client and one or more remote hosts to build a system called DYNAMINER [8]. The system leverages the "conversation" aspect of web sessions to build *web conversation graphs*, study their graph-centric properties, and build a web conversation classifier used for real-time detection. The key insight in DYNAMINER is the capturing of crucial sub-behaviors of redirection chains, malware download, and post-infection dynamics in web-borne malware infection. From live web sessions, DYNAMINER is guided by pre-learned threat clues to build conversation graphs from which it extracts graph-centric properties to classify the graphs as infection or benign. On forensic and live case studies, DYNAMINER outperformed the state of the art by detecting previously unseen malware 11 days earlier.

Cyber-crime Analysis: Turning the Table Against Cybercriminals

Counter-Offensive Infiltration of Exploit Kits: This research was motivated by an observation that revealed vulnerabilities in exploit kit source code and the need to explore if one can use vulnerabilities in exploit kits to infiltrate them and turn the table against the cybercriminals behind exploit kits. To this end, we developed a toolkit, called EKHUNTER [5], that enables cyber-crime analysts to conduct counter-offensive infiltration on exploit kit servers. EKHUNTER creatively combines multi-faceted vulnerability analysis of exploit kit source code with constraint solving to generate working exploits that can be issued as HTTP requests to active exploit kits on the Web. This toolkit is the first of its kind in the counter-offensive front of fighting cyber-crime on the Internet. This work practically demonstrated the feasibility of using the very exploitive techniques of cybercriminals against their own malicious toolkits to reverse-fingerprint attack toolkits and stop their malicious deeds. In addition, I have demonstrated using real malicious toolkits that EKHUNTER can be used by ethical cyber threat analysts, law enforcement, or cybersecurity vendors for the purpose of cyber-crime analysis and (active and passive) cyber-crime investigation.

Advanced and Persistent Threats: Needle-in-a-Haystack

Advanced and Persistent Threats (APTs) have hit close to home in many sectors. It has become customary to witness APTs in a number of high profile cases such as customer data breaches, government data theft, political organization hacks, and infiltration of critical infrastructures. APTs

often combine social engineering (e.g., spear phishing) with advanced exploits to bypass first-line defenses such as Address Space Layout Randomization and sandboxes. As a result, enterprises resort to second-line defenses such as Intrusion Detection Systems and Security Incident and Event Management. Nevertheless, we continue to witness many APT campaigns slip through the radar of second-line defenses and remain undetected for months.

The key challenges in the detection and analysis of APTs are (i) the "needle in a haystack" phenomenon —spotting the tiny fraction of real attacks from an ocean of data emitted by enterprise systems; and (ii) "connecting the dots" —how to stitch isolated steps together toward an actionable attack story that shows the entry points, details of attack activities, and exit points.

To address these challenges, we developed SLEUTH [13], a system that detects attacks in real-time and constructs an attack campaign graph that summarizes an attack campaign, with low memory footprint and fast analytics. SLEUTH addresses attack detection and campaign reconstruction by tracking and reasoning about dependencies between entities such as files, processes, and network sockets, abstracted in a *provenance graph*. For attack detection, SLEUTH relies on *provenance-based policies* that capture benign and suspicious uses of (a) data and code from untrusted sources, and (b) confidential data. SLEUTH was evaluated on an adversarial engagement led by a dedicated red team from a government agency. The red team launched APT-like campaigns against target machines over a period of two weeks. In the same period, SLEUTH received the event streams of the target machines and was able to spot and correctly reconstruct attack campaigns from millions of events with about 99.9% benign activities. For instance, on a campaign of 38.5M events, SLEUTH produced an attack campaign graph of only 130 events, with 5 orders of magnitude reduction in event volume.

Web Security: Vulnerability Analysis and Automated Exploit Generation

Web Application Configuration Vulnerability: Although security misconfiguration is among the top-10 web application security risks over the years, it received little attention in the state-of-the-art vulnerability scanners. To fill this gap, I developed CONFEAGLE [11], an automated approach to analyze and detect configuration vulnerabilities in web applications. The approach combines web application structure scanning, lightweight code analysis, and established vulnerability scoring metrics to automatically assess the configuration security posture of web applications. CONFEAGLE uncovered serious configuration vulnerabilities in 14 widely deployed open source web applications and those vulnerabilities were not detected by other web application vulnerability scanners at the time. The vulnerabilities that CONFEAGLE discovered are very serious as the web applications are downloaded and deployed in millions. Moreover, those web applications are the backbones of many critical web services in sectors such as healthcare, education, and e-commerce. CONFEAGLE has been open-sourced since 2010 and has been in use in over 50 countries.

Automated Exploit Generation: For quite a while, vulnerabilities reported by state of the art white-box web application vulnerability analysis techniques suffered from high false positives. The main challenge is the difficulty in converting vulnerabilities into concrete exploits. To address this challenge, we developed CHAINSAW [1], a system that automatically constructs a sequence of malicious HTTP request inputs from a web application source code so as to direct an application's execution to an exploitable sink. Unlike exploit generation for standard binaries, CHAINSAW has to tackle challenges associated with typical web application characteristics: their multi-module nature, interposed user input, and multi-tier architectures that involve databases with constraints. We combined static code analysis with SAT solvers to generate concrete exploits. CHAINSAW was tested on 9 open source applications and generated over 199 first- and second-order injection exploits combined, significantly outperforming the state of the art.

Pervasive Computing: Context-Awareness and Anomaly Detection

Context-Awareness: For my master's thesis, I addressed lack of consideration for quality of context information and limitation of existing approaches to incorporate domain requirements in pervasive

healthcare. Using an ontology-based representation of domain concepts in pervasive healthcare and through a novel context information refinement engine that considers service parameters set by a physician on a mobile device, I developed and evaluated a context information refinement framework [6] that tailors information to the current context (e.g., location, activity) and service parameter specifications of a physician (e.g., "don't notify me when I am in the OR").
Anomaly Detection: In the same domain, I also collaborated on a host-based anomaly detection problem that demonstrated how to model nominal activities of physicians by monitoring their day-to-day routine on a mobile device. Using rules mined from nominal activities, we developed a system [2] that can detect deviations from intended behavior and flag anomalies such as attempts to escalate privileges.

Future Research Directions

As more and more cyber services and systems get connected to the Internet, they get exposed to remote manipulation by cyber-criminals ranging from script kiddies to organized groups/states whose motivation is often economic or political. I believe the future of cyber threats defense will highly depend on scalability, reliability, and evolution-awareness of threat-centric analytics driven by big-data generated from heterogeneous systems and devices. To this end, I envision my future research along three lines: data-driven security analytics on enterprise systems, IOC-driven threat detection and forensics, and large-scale cyber threat hunting.

Threat-Centric Event Stream Analytics for Enterprises

As cyber-attacks become more stealthy, auditing systems at a level of granularity that can assist cyber threat detection and forensics has become an important research focus in recent years. In line with my interest in using machine learning and scalable big-data analytics techniques for cybersecurity, in the near future, I plan to focus on developing the foundational techniques and tools for a real-time OS event analytics system that assists system admins and cyber threat analysts in making informed decisions on observed or potential cyber threats. The classic challenges of big-data analytics including volume, velocity, variety, data provenance, granular auditing, and efficient and scalable analytics and visualization, are the research challenges that I plan to tackle by leveraging my experience and the expertise of colleagues working on topics of synergy with my research visions.
Toward this goal, I am currently investigating OSQUERY (`https://osquery.io/`), a cross-platform auditing engine that exposes the OS activities as a huge database. The exposure of the whole OS as a huge database is an opportunity and a challenge at the same time. It is an opportunity because it allows cyber threat analysts to craft and execute threat-centric queries both in real-time and in a forensic setting. It poses challenges with regards to (a) configuring the auditing, (b) scheduling of queries, (c) analytics algorithms, and (d) visual presentation of detected cyber-threats for threat analysts. In the long term, I will carefully explore techniques and build systems that not only address these research challenges, but also result in production systems with real-life impacts.

IOC-Driven Compromise Detection and Forensics

To detect and investigate intrusion, the cybersecurity industry relies on Indicator of Compromise (IOC) as forensic artifacts observed in network traffic or operating system logs. However, beyond automating the extraction of IOCs (e.g., IPs, domains, and hashes of malware)from open-source cyber threat intelligence sources (e.g., security blogs), little or no attention is paid to the enrichment of IOCs to make them of vital value for (a) real-time compromise detection and (b) forensic analysis. I plan to advance the state of the art/practice with regards to (a) and (b).
My short-term goal in this space is to develop techniques toward semantically rich and evasion resilient IOCs by combining multiple threat intelligence sources so as to maximize the quality and quantity of IOC knowledge-base. In the long run, I aim to explore the practical viability and the challenge thereof of field-testing automatically extracted IOCs for detection and forensics of cyber

threats. In this line of work, I plan to tailor NLP techniques for the domain-specific requirements of IOC extraction from natural language sources, and big-data analytics for engineering effective compromise detection and impact analysis.

Toward Precise Infection Clue Inference for Large-Scale Cyber Threat Hunting

In my recent work DYNAMINER [8], the experimental evaluation revealed that, for real-time detection and on-the-fly characterization, it is generally challenging to initiate threat alarms due to the predominantly benign web transactions interleaved with malicious transactions. As threats evolve, a detection strategy that is powered by a static set of heuristics for threat clue inference will be outdated. As a result, a detection system built as such will miss real clues (resulting in false negatives) or include non-clues (resulting in false positives) in its threat abstraction. For large-scale threat hunting, it is important to continuously tune threat clue models to the changes in the threat landscape. My research plan is to learn representative models for such special cases in order to cut false signals significantly by leveraging a model pool of anticipated false signals. To save computing resources and cyber analysts' time, one strategy to explore is to learn these models a priori and then query them as a sanity check just before putting a label on a web conversation graph. The challenge then is how to keep up with the evolving threat landscape. Again, to address this, I envision a carefully crafted, automated update strategy with room for incorporating exceptions such as compromised software repositories and executables with backdoors.

References

[1] Abeer Alhuzali, Birhanu Eshete, Rigel Gjomemo, and V.N. Venkatakrishnan. Chainsaw: Chained automated workflow-based exploit generation. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, pages 641–652, 2016.

[2] Biniyam Asfaw, Dawit Bekele, Birhanu Eshete, Adolfo Villafiorita, and Komminist Weldemariam. Host-based anomaly detection for pervasive medical systems. In *CRiSIS 2010, Proceedings of the Fifth International Conference on Risks and Security of Internet and Systems*, pages 1–8, 2010.

[3] Birhanu Eshete. *Effective Analysis, Characterization, and Detection of Malicious Activities on the Web*. PhD thesis, University of Trento, Italy, 2013.

[4] Birhanu Eshete. Effective analysis, characterization, and detection of malicious web pages. In *22nd International World Wide Web Conference, WWW, Companion Volume*, pages 355–360, 2013.

[5] Birhanu Eshete, Abeer Alhuzali, Maliheh Monshizadeh, Phillip A. Porras, Venkat N. Venkatakrishnan, and Vinod Yegneswaran. Ekhunter: A counter-offensive toolkit for exploit kit infiltration. In *22nd Annual Network and Distributed System Security Symposium, NDSS*, 2015.

[6] Birhanu Eshete, Dawit Bekele, Adolfo Villafiorita, and Komminist Weldemariam. Context information refinement for pervasive medical systems. In *The Fourth International Conference on Digital Society, ICDS*, pages 210–215, 2010.

[7] Birhanu Eshete and V. N. Venkatakrishnan. Webwinnow: Leveraging exploit kit workflows to detect malicious urls. In *Proceedings of the 4th ACM Conference on Data and Application Security and Privacy*, CODASPY '14, pages 305–312, 2014.

[8] Birhanu Eshete and V. N. Venkatakrishnan. Dynaminer: Leveraging offline infection analytics for on-the-wire malware detection. In *47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN*, pages 463–474, 2017.

[9] Birhanu Eshete, Adolfo Villafiorita, and Komminist Weldemariam. Early detection of security misconfiguration vulnerabilities in web applications. In *Sixth International Conference on Availability, Reliability and Security, ARES*, pages 169–174, 2011.

[10] Birhanu Eshete, Adolfo Villafiorita, and Komminist Weldemariam. BINSPECT: holistic analysis and detection of malicious web pages. In *Security and Privacy in Communication Networks - 8th International ICST*, pages 149–166, 2012.

[11] Birhanu Eshete, Adolfo Villafiorita, Komminist Weldemariam, and Mohammad Zulkernine. Confeagle: Automated analysis of configuration vulnerabilities in web applications. In *IEEE 7th International Conference on Software Security and Reliability,SERE*, pages 188–197, 2013.

[12] Birhanu Eshete, Adolfo Villafiorita, Komminist Weldemariam, and Mohammad Zulkernine. EINSPECT: evolution-guided analysis and detection of malicious web pages. In *37th Annual IEEE Computer Software and Applications Conference, COMPSAC*, pages 375–380, 2013.

[13] Md Nahid Hossain, Sadegh M. Milajerdi, Junao Wang, Birhanu Eshete, Rigel Gjomemo, R. Sekar, Scott Stoller, and V. N. Venkatakrishnan. SLEUTH: real-time attack scenario reconstruction from COTS audit data. In *26th USENIX Security Symposium, USENIX Security*, pages 487–504, 2017.