NJIT

Department of Computer Science

Securing Operating System Kernels with Fewer Shots

Yueqi Chen
Ph.D. candidate at Pennsylvania State University and a Predoctoral Fellow at Northwestern University

Hosted by Vincent Oria

Date: Wednesday February 23, 2022
Coffee: 10:45 AM – 11:00 AM
Time: 11:00 AM – 12:00 PM
Location: GITC 4402 (4th Floor Seminar Lecture Hall)
WebEx Link: **https://njit.webex.com/njit/j.php?MTID=m77f7acc3ba5934d04c7e1fd552270225**

https://cs.njit.edu/seminars

Abstract: Despite significant efforts on cybersecurity, we are observing more and more attacks in recent years. The reason behind this is all our efforts aim at individual incidents and there is no deep understanding of attack surfaces in software systems. As a result, software systems are integrated with too many individual patches and ad-hoc protections (just like one shot per variant), which brings painful overhead but not substantial security benefits.

In this talk, I will present a systematic approach to understanding attack surfaces. This approach provides security analysts and developers with the ability to quantify the impact of attack surfaces and facilitate the development of universal and effective defense solutions. Technically, this approach consists of two steps - induction and deduction. The induction step abstracts the essential causality behind the individual attack incident and the deduction step applies the abstraction to different contexts (e.g., various systems independently created by different vendors). In this talk, I will exemplify this induction and deduction approach with a security incident in the Linux kernel. Following this, I will present a general and effective defense that mitigates the induced attack surfaces and is widely adopted in various commodity Operating System kernels.

In the future, I plan to further advance this systematic approach and make it a fundamental part of the entire software development lifecycles. More specifically, I will: 1) enrich induction and deduction techniques for more attack forms under new contexts, 2) improve the scalability of these techniques via automation, and 3) optimize and re-construct existing defenses to build a new and comprehensive architecture that mitigate attack surfaces in a quantitative approach.

Bio:

Yueqi Chen received his B.Sc degree from Nanjing University. He is currently a fifth-year Ph.D. student at Pennsylvania State University and a predoctoral fellow with Dr. Xinyu Xing at Northwestern University. He was awarded the IBM Ph.D. Fellowship 2020-2022. In general, his research focuses on system security and software security. He is particularly interested in developing systematic approaches to inducing, deducting, and mitigating attack surfaces. Along this thread, he has published 10 papers in top-tier academic conferences, including IEEE S&P, USENIX Security, ACM CCS, NDSS, OOPSLA, ACM/IEEE ICSE, IEEE/ACM ASE as leading authors and co-authors over the past 4 years. In addition, he presented his works at CLK 2019, BlackHat Europe 2019, BlueHat IL 2020, LSS Europe 2020, BlackHat Aisa 2021, LSS North America 2021, BlackHat Europe 2021. His research works were covered by high-profile media (e.g., Dark Reading) and have received wide recognition from the industry, including Amazon, Apple, Baidu, Google, Grsecurity, IBM, JD.com, Linux Foundation, Microsoft, and Red Hat. His work is integrated into the internal threat alerting platform of JD.com. The new defenses in AutoSlab produced by Grsecurity and iPhone 13 series from Apple can find their prototypes in his works. As a team member of r3kapig and Nu1L, he participated in DEF CON 26 CTF Final and DECONF 29 CTF Final, and ranked 16th and 7th, respectively. He ranked 5th in NSA codebreaker 2017.