ThreatSCOPE

Addressing software vulnerability in embedded systems

The team at cyber-security firm BlueRISC Inc., is led by Dr Csaba Andras Moritz, Founder and Chairman of BlueRISC Inc. and Professor of Electrical and Computer Engineering at the University of Massachusetts Amherst. They have developed ThreatSCOPE, a system assurance tool that analyses and mitigates software vulnerabilities and cyber threats in embedded systems. ThreatSCOPE uses BlueRISC's patented ImmuneSoft technology to offer both static and runtime cyber-hardening.

Embedded systems controlling complex systems, including medical devices, avionics, automotive and the Internet of Things (IoT), are becoming increasingly sophisticated. The software supporting these systems is growing more and more complex. As this complexity increases, so does the challenge in identifying vulnerabilities within the software, thus exacerbating the likelihood of cyber threats. These embedded systems rely on external data, such as user input and external sensors, and may be connected to public networks. Attackers can exploit the unique interactions that take place between systems and between users and systems in order to find weaknesses that enable them to leak information or take over control of the system. Moreover, it is impossible to test for these complex, data driven interactions. This makes the evaluation of a system's cyber-security attack surface a challenging problem.

The research and development team at cyber-security firm BlueRISC Inc., led by Dr Csaba Andras Moritz, entrepreneur and Professor of Electrical and Computer Engineering at the University of Massachusetts Amherst, have developed ThreatSCOPE to address this software vulnerability problem. ThreatSCOPE is a system assurance tool that analyses software vulnerabilities and related cyber threats. Professor Moritz describes how ThreatSCOPE builds on novel binary analysis concepts in order to identify where exploitable code exists within an embedded system. It then determines the paths that lead to potential exploitation of the system. The graphical front-end enables users to visualise the exploitability analysis results together with a breakdown of any potential vulnerabilities that have been discovered. ThreatSCOPE can also generate automated reports, presenting the results of its analyses at various levels including an executive summary.

Pavel Ignatov/Shutterstock.com

BINARY-LEVEL FRAMEWORK

ThreatSCOPE is underpinned by a binary-level framework that supports a wide range of embedded CPUs (central processing units) and environments and facilitates vulnerability analysis in final binaries in embedded systems. It can automatically extract artifacts (i.e., program analytic information that represent essential aspects of what the program does) that contribute to weaknesses in the system and identifies potentially exploitable paths from both application and system software. The R&D team have designed ThreatSCOPE to perform automatic binary/executable reverse-engineering of the system of interest, as well as extracting and characterising possible security weaknesses.

EXPLOITABILITY

Instead of depending on prior knowledge of existing vulnerabilities and solutions, ThreatSCOPE's exploitability characterisation technology extracts the conceptual artifacts that make up the operational requirements of what would be a successful cyber-attack on the embedded system. Dr Moritz explains that essentially any attempts to exploit an embedded system has to involve either system-to-system or user-to-system interactions at some point. Normally these interactions manage the system's underlying functionality. If, however, these interactions have been created or formatted in a particular way they can offer a path for hackers to gain access in order to leak information from it or take control and exploit it.

Visual Generation /Shutterstock.com

ThreatSCOPE builds on novel binary analysis concepts in order to identify where exploitable code exists within an embedded system

The firmware's underlying functionality is also of special interest to the researchers, as only some types of codes actually meet the criteria for exploitation. A potentially exploitable path occurs when these artifacts, or codes meeting the exploitability requirements, are associated with locations where interactions with external interfaces take place. Once the exploitability artifacts have been identified, ThreatSCOPE generates an Exploitability Artifact Graph (EAG) using the execution paths to establish the logical connections between the artifacts. The EAG offers users an interactive view of the procedures that contain the specific functionality and connectivity that satisfy the exploitability criteria. It provides a visualization of the embedded system software's attack surface, identifying procedures with external interfaces that communicate with either users or another information processing system, procedures with potentially modifying vulnerabilities and those with weaknesses that could leak data from the system.

REVERSE-ENGINEERING

When an embedded device's firmware image or software executable file is input, the ThreatSCOPE toolkit performs an automated reverse engineering to find out how the firmware or executable actually works in terms of its underlying functionality, procedures and communication. This provides an intermediate representation of the input that is architecture agnostic so the analyses can be performed symbolically without the need for any knowledge of the underlying architecture.

STATIC VULNERABILITY CHARACTERIZATION

The toolkit then performs automated program analyses based on this intermediate representation, in order to identify and characterize any exploitability artifacts that could lead to potential security weaknesses. ThreatSCOPE characterizes these artifacts in relation to either the system's data-flow in relation to the external input or output interfaces, or the type of attack and whether they could enable modifying or leak data from the system. Once categorised, these artifacts together with the execution paths that connect them make up the exploitation-relevant locations shown in the EAG. ThreatSCOPE captures a superset of

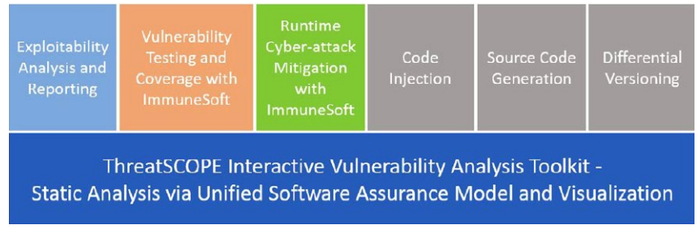

ThreatSCOPE static analysis and runtime capabilities.

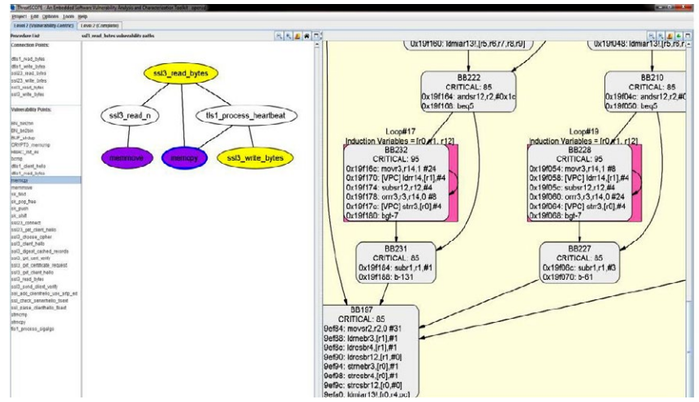

ThreatSCOPE visualization showing multiple granularities of viewing existing security defects and the associated artifacts.

Visual Generation /Shutterstock.com

vulnerable codes from the software via the exploitability artifacts and their interactions. While some of these potentially exploitable codes can be fixed statically, others depend on particular runtime conditions that are not traceable during testing. Dr Moritz emphasises that it is essential to enable the embedded software to detect the occurrences of these exploitations.

IMMUNESOFT – A CYBER IMMUNE SYSTEM

ThreatSCOPE's static analysis also allows the addition of a runtime component, a self-contained software module, to the embedded systems in order to detect and possibly provide software healing for the application should an attack be attempted. BlueRISC have developed and patented ImmuneSoft technology which uses information that is extracted statically to identify and harden possible weaknesses, in order to defend and secure the system by reducing its attack surface. ImmuneSoft can be employed during both vulnerability testing and runtime cyber-attack mitigation.

VULNERABILITY TESTING

During the vulnerability testing phase, ImmuneSoft codes are generated automatically and slotted into the binary code transparently where either artifacts or their paths have been identified. These codes can be used to generate vulnerability-centric coverage metrics during both functional and penetration testing that inform the user how well these potentially vulnerable codes have actually been tested. This information is not available elsewhere. ImmuneSoft codes can also identify the conditions that enabled the exploitation of a particular artifact or path to succeed.

ThreatSCOPE's ImmuneSoft technology is the first system to target healing of an embedded software during a cyber-attack

RUNTIME CYBER-ATTACK MITIGATION

ImmuneSoft codes can also remain in the deployed application and perform runtime cyber-attack mitigation. These codes provide a cyber-hardening solution with low overheads that can detect and respond to runtime cyber-attacks. In cases where the user has the required knowledge of the application/interface, they can develop a bespoke ImmuneSoft code for that particular application/interface that ThreatSCOPE can automatically insert into the embedded software. Alternatively, ThreatSCOPE is able to automatically generate ImmuneSoft codes which can identify the conditions that are required for a successful attack to be carried out on a vulnerability artifact during runtime. When a cyber-attack is detected at runtime, the ImmuneSoft codes activate a response mechanism. Possible response mechanisms range from those, such as logging, that have no active impact on the application to more aggressive responses such as autonomous software healing.

FUTURE DEVELOPMENTS

Exploitability in networked embedded systems is a growing problem and it is no longer sufficient to just test for known threats. ThreatSCOPE and ImmuneSoft offer both static and runtime cyber-hardening solutions. The team also uses AI extensively to support source-level analysis in addition to executables. These technologies can be used during a system's development process for analysis and testing, as well as in deployed systems. Dr Moritz and his colleagues at BlueRISC are currently developing a range of additional features that incorporate support for instrumentation, code-injection and modelling into the ThreatSCOPE toolkit.

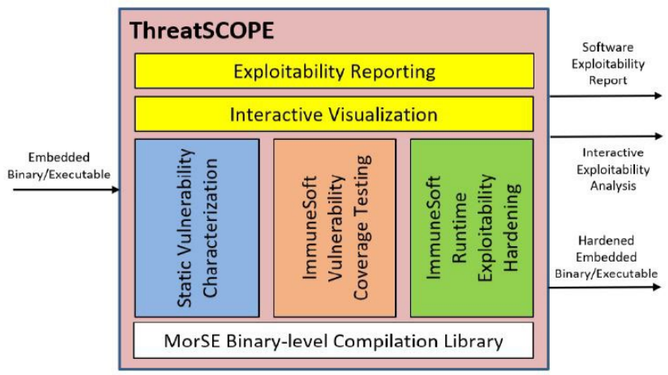

The ThreatSCOPE and ImmuneSoft toolkit architecture.

Behind the Research



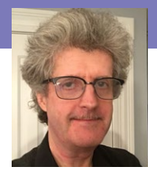

Dr Csaba Andras Moritz | Kristopher Carver | Karl Zimmerman | Jeff Gummeson | Carl Seneca

E: andras@bluerisc.com **T:** +1 (617)-517 6324 **W:** www.bluerisc.com

Detail

Address: BlueRISC Inc.,
400 Amity Street, Suites 0-1-3-4
Amherst, 01002 MA. USA

Bio

Dr Csaba Andras Moritz is the founder and Chairman of the cyber-security firm BlueRISC Inc. He is the inventor and architect of BlueRISC's R&D directions, including related to its hardware security offerings, system assurance tools, post-compromise analysis, and EPRIVO for digital privacy. He is also a Professor of Electrical and Computer Engineering at University of Massachusetts Amherst.

Kristopher Carver is the Technical Director at BlueRISC and the PI on many cyber-security R&D projects with both DoD and DHS. He has talked in leading cyber-security conferences and lead the R&D of several system-assurance and embedded security products developed at BlueRISC.

Karl Zimmerman is a Senior Software Engineer. He is the lead compiler developer behind the ThreatSCOPE tool and is an expert on static analysis, binary level analysis, and cyber security threats. A graduate from MIT, he worked in several companies in various engineering roles before assuming a leading role in BlueRISC system assurance tools and technologies area.

Jeff Gummeson is a Senior Security Architect at BlueRISC and an expert on several cyber-security directions BlueRISC researches. He is a PI on multiple R&D projects with DHS and DoD, a speaker at several cyber security conferences, and leads BlueRISC's post-compromise analysis research and WindowsSCOPE product line.

Carl Seneca is a Senior Software Engineer at BlueRISC. He is a lead investigator of cyber threats as they pertain to the ThreatSCOPE technology including MITRE's Common Weakness Enumeration frameworks and the National Vulnerability Database. Carl is also an expert on protection approaches with secure processors.

Funding
DHS, ONR, US Air Force, DARPA, US Army.

Research Objectives

BlueRISC Inc. develops tools and systems to tackle issues as wide-ranging as vulnerability analysis/avoidance to exploit detection and software self-healing in the field, hardware-assisted protections, and post cyber-attack forensic analysis. It has completed advanced cyber-security R&D and been sponsored by several research organisations in the US including NSF, DARPA, DHS, and others. BlueRISC Inc. has customers in 20 countries.

References

Moritz, C.A., Carver K, Zimmerman K., Gummeson J., Seneca C., (2020). White Paper: ThreatSCOPE Exploitability Analysis and Mitigation. [online] BlueRISC. Available at: https://www. bluerisc.com/whitepapers/ [Accessed 23rd November 2020]

Patent 941196. Characterizing, Detecting and Healing Vulnerabilities in Computer Code. Moritz, C.A., Carver K, Gummeson J.

Patent 9754112. Detection and Healing of Vulnerabilities in Computer Code. Moritz, C.A., Carver K, Gummeson J.

Personal Response

What is the most rewarding outcome from ThreatSCOPE to date?

"ThreatSCOPE initially started as a research project funded by DHS and DARPA but has since evolved into a flagship approach for binary-level vulnerability analysis of embedded systems. It follows an unconventional mindset: rather than simply looking for known threats, ThreatSCOPE tries to identify the fundamental reasons why a software is vulnerable. That enables detecting zero-day vulnerabilities. ThreatSCOPE has been ranked as the highest performing tool by a world-leading embedded system vendor and is increasingly targeted to a variety of markets from defence systems, to automotive, industrial control, and medical devices. Due to its fundamental nature, it has rapidly evolved and has also become a precursor to BlueRISC's hardware RTL and IC analysis technology that performs similar cyber weakness analysis on hardware designs towards finding malicious circuits and vulnerabilities."