POINTERS

Wyvern programming language builds secure apps

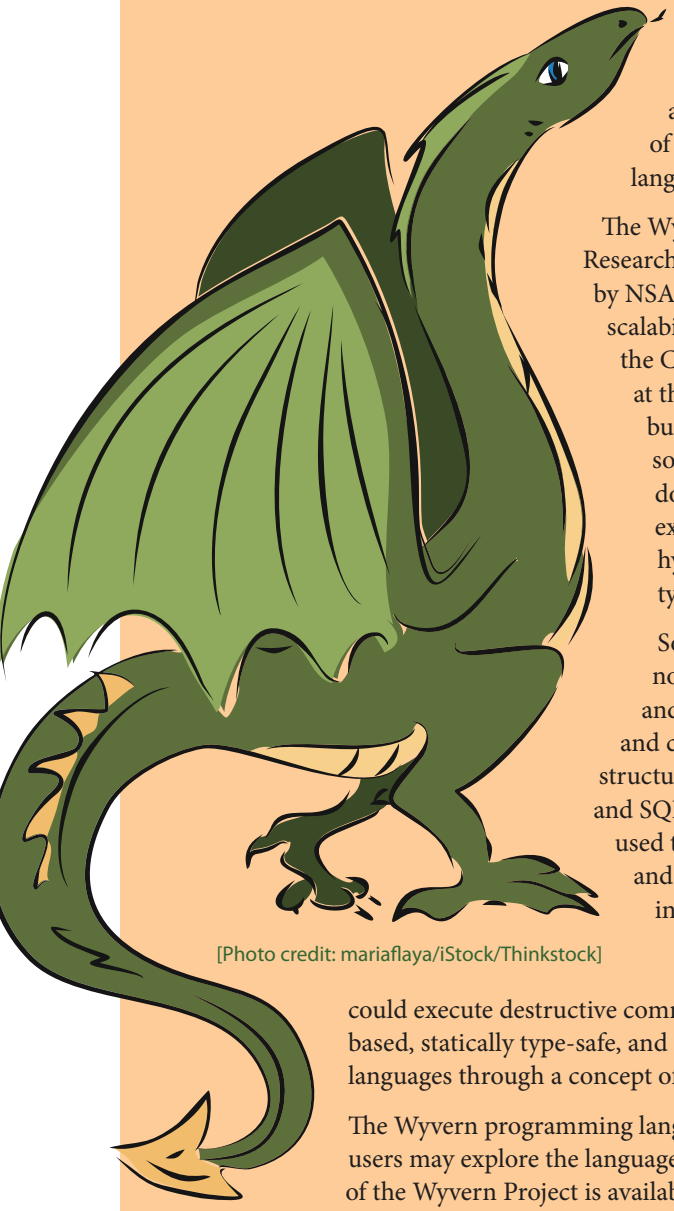

[Photo credit: mariaflaya/iStock/Thinkstock]

In legend, the dragon-like creature known as the wyvern used its fiery breath and poisonous bite to protect its treasure from thieves. Still a popular figure on coats of arms and in electronic games, the wyvern has also inspired researchers at the Carnegie Mellon University (CMU) Science of Security (SoS) Lablet to develop the Wyvern extensible programming language to keep web and mobile applications more secure.

The Wyvern project, led by Dr. Jonathan Aldrich of the Institute for Software Research in the CMU School of Computer Science, is part of research supported by NSA and other agencies to address hard problems in cybersecurity, including scalability and composability. Dr. Aldrich, in collaboration with researchers at the CMU SoS Lablet, New Zealand colleague Dr. Alex Potanin, and researchers at the Victoria University of Wellington, have been developing Wyvern to build secure web and mobile applications. The language is designed to help software engineers build those secure applications using several type-based, domain-specific languages within the same program. Wyvern is able to exploit knowledge of sublanguages (e.g., structured query language (SQL), hypertext markup language (HTML), etc.) used in the program based on types and their context, which indicate the format and typing of the data.

Software development has come a long way, but the web and mobile arenas nonetheless struggle to cobble together different languages, file formats, and technologies. This proliferation is inefficient and thwarts scalability and composability. For example, a typical web page might require HTML for structure, cascading style sheets for design, JavaScript to handle user interaction, and SQL to access the database back-end. The diversity of languages and tools used to create an application increases the associated development time, cost, and security risks. It also creates openings for cross-site scripting and SQL injection attacks. Wyvern eliminates the need to use character strings as commands, as is the case, for instance, with SQL. By allowing character strings, malicious users with a rough knowledge of a system's structure could execute destructive commands. Instead, Wyvern is a pure object-oriented language that is value-based, statically type-safe, and supports functional programming. It supports HTML, SQL, and other web languages through a concept of composable type-specific languages.

The Wyvern programming language is hosted at the open-source site GitHub; interested potential users may explore the language at https://github.com/wyvernlang/wyvern. In addition, a description of the Wyvern Project is available on the Cyber Physical Systems Virtual Organization web page at http://cps-vo.org/node/21424.

Building secure and resilient software from the start

Researchers at the North Carolina State University (NCSU) SoS Lablet are addressing the hard problems of resilience and predictive metrics in the development of secure software. The team, under principal investigators Laurie Williams and Mladen Vouk, empirically determined that security vulnerabilities (i.e., faults that violate implicit or explicit security policies) can be introduced into a system because the code is too complex, changed too often, or not changed appropriately. According to one NCSU study, source code files changed by nine developers or more were 16 times more likely to have at least one vulnerability after release. From such analyses, researchers can develop and disseminate predictive models and useful statistical associations to guide the development of secure software.

NCSU researchers have identified two key questions whose answers will require rigorous analysis and testing:

1. If highly attack-resilient components and appropriate attack sensors are developed, will it become possible to compose a resilient system from these component parts?
2. If so, how does that system scale and age?

One very simple and effective defensive strategy would be to build in a dynamic application firewall that does not respond to "odd" or out-of-norm inputs, such as those associated with zero-day attacks. While not foolproof, this approach would force attackers to operate within an application's "normal" operational profile.

The research has generated tangible results. Three recent papers have been written as a result of this research. "On coverage-based attack profiles," by Anthony Rivers, Mladen Vouk, and Laurie Williams (doi: 10.1109/SERE-C.2014.15); "A survey of common security vulnerabilities and corresponding countermeasures for SaaS," by Donhoon Kim, Vouk, and Williams (doi: 10.1109/GLOCOMW.2014.7063386); and "Diversity-based detection of security anomalies" by Roopak Venkatakrishnan and Vouk (doi: 10.1145/2600176.2600205). For the complete SoS newsletter article, visit http://www.cps-vo.org/node/21426.

Interest in cybersecurity science heats up at HotSoS 2015

The 2015 Symposium and Bootcamp on the Science of Security (HotSoS) took place April 21–22 at the UIUC National Center for Supercomputing Applications. This third annual conference, part of the NSA SoS project, brought together researchers from numerous disciplines seeking a rigorous scientific approach toward cyber threats. David Nicol, director of the Information Trust Institute and coprincipal investigator for the UIUC SoS Lablet, was conference chair. Kathy Bogner, Intelligence Community coordinator for cybersecurity research, represented the NSA sponsor. Featured speakers included Mike Reiter, Lawrence M. Slifkin distinguished professor of computer science, University of North Carolina; Jonathan Spring, researcher and analyst for the Computer Emergency Response Team division of the Software Engineering Institute, CMU; and Patrick McDaniel, professor of computer science and director of the Systems and Internet Infrastructure Security Laboratory, Penn State University.

Five tutorials and a workshop took place concurrently with individual presentations. Tutorials covered

[Photo credit: Sifis Diamantidis/Hemera/Thinkstock]

social network analysis; human behavior; policy-governed secure collaboration, security-metrics-driven evaluation, design, development and deployment; and resilient architectures. Thirteen individual papers, presented by British and US researchers, covered the following: signal intelligence analyst tasks, detecting abnormal user behavior, tracing cyber-attack analysis processes, vulnerability prediction models, preemptive intrusion detection, enabling forensics, global malware encounters, workflow resiliency, sanctions, password policies, resource-bounded systems integrity assurance, active cyber defense, and science of trust. The agenda and presentations are available from the Cyber-Physical Systems Virtual Organization (CPS-VO) website. Members' may access the site at http://cps-vo.org/node/3485/browser. Nonmembers may access the site at http://cps-vo.org/group/SoS.

Next year's event will take place in Pittsburgh and be hosted by the CMU SoS Lablet.

Adoption of Cybersecurity Technology Workshop

The Special Cyber Operations Research and Engineering (SCORE) subcommittee sponsored the 2015 Adoption of Cybersecurity Technology (ACT) Workshop at the Sandia National Laboratories in Albuquerque, New Mexico, from 3–5 March 2015. The vision for the workshop was to change business practices for adoption of cybersecurity technologies; expose developers, decision makers, and implementers to others' perspectives; address the technology, process, and usability roadblocks to adoption; and build a community of interest to engage regularly.

Attendees represented four segments of the cybersecurity world: researchers and developers, decision makers, implementers, and experts on human behavior. They explored, developed, and implemented action plans for four use cases that addressed each of the four fundamental cybersecurity goals: 1) device integrity, 2) authentication and credential protection/defense of accounts, 3) damage containment, and 4) secure and available transport. This construct provided a framework specifically designed to confront spear phishing.

Participants were primarily government personnel, with some individuals from federally funded research and development centers, academia, and industry. These cybersecurity professionals believe that as many as 80% of their field's current problems have known solutions that have not been implemented.

The workshop is the kickoff activity for a sustained effort to implement cybersecurity technology solutions throughout the US government, and it is expected to become an annual event.

The agenda focused on specific threat scenarios, cohorts' concerns to promote understanding among groups, addressing the use cases, and developing action plans for implementation via 90-day phases known as "spins." Participants will brief the spin results to the ACT organizing committee and to threat analysts who will assess progress.

In order to illuminate systemic barriers to adoption of security measures, the workshop focused specifically on countering the threat from spear phishing, a social-engineering trick that hackers use to convince people to provide passwords, financial data, and other private information.

The goal over the next year is to strengthen government networks against spear phishing attacks by applying the selected technologies identified in the four use cases. Although this activity is tactical in nature, it provides an underlying strategy for achieving broader objectives as well as a foundation for transitioning collaborative cybersecurity engagements.

For the complete article, visit: http://www.cps-vo.org/node/21405.

Sandboxing in the cloud

Cybersecurity experts working with complex systems often have to isolate certain components that cannot be fully trusted. A common technique for this is known as sandboxing—a method in which security layers are put in place to encapsulate software components and impose policies that regulate interactions between the isolated components and the rest of the system. Sandboxes provide a mechanism to work with a full sampling of components and applications, even when they are known to be malicious. Of course, if the sandbox fails or is bypassed, the production environment may be compromised.

Michael Maass, William L. Scherlis, and Jonathan Aldrich from the Carnegie Mellon University (CMU) Institute for Software Research in the School of Computer Science propose a cloud-based sandbox method to mitigate the risk of sandbox failure and raise the bar for attackers. The researchers call their approach "in-nimbo sandboxing," after *nimbus,* the Latin word for "cloud." Using a technique they liken to software-as-a-service, they tailor the sandbox to the specific application in order to encapsulate components with smaller, more defensible attack surfaces than other techniques. This remote encapsulation reduces the likelihood of a successful attack as well as the magnitude or degree of damage from a successful attack.

The authors note that "most mainstream sandboxing techniques are in-situ, meaning they impose security policies using only trusted computing bases within the system being defended. Existing in-situ sandboxing approaches decrease the risk that a vulnerability will be successfully exploited because they force the attacker to chain multiple vulnerabilities together or bypass the sandbox. Unfortunately, in practice, these techniques still leave a significant attack surface, leading to a number of attacks that succeed in defeating the sandbox."

With funding from the CMU Science of Security (SoS) Lablet, Maass, Scherlis, and Aldrich conducted a field trial with a major aerospace firm. They were able to compare an encapsulated component deployed in an enterprise-managed cloud, with the original, unencapsulated component deployed in the relatively higher-value user environment. The researchers' assessment was based on the criteria of performance, usability, and security.

1. **Performance evaluation:** The trial focused on latency of interactions and ignored resource consumption. For deployed applications, the technique increased user-perceived latency only slightly.
2. **Usability evaluation:** The sandbox's design was structured to present an experience identical to the local version, and users judged that this was accomplished. Results suggest that the in-nimbo approach may be feasible for other types of systems as well because a field trial system is built primarily using widely adopted established components.
3. **Security evaluation:** Cloud-based sandboxes are more difficult to attack, partly because defenders can customize the environment in which an encapsulated computation takes place and partly because of the inherently ephemeral nature of cloud-computing environments. The in-nimbo system separates a component of interest from its operating environment and replaces it with a specialized transduction mechanism to manage interactions with the now-remote component, which has been moved to the cloud environment.

The authors indicate that this work is a precursor to an extensive study, still in progress, that evaluates more than 70 examples of sandbox designs and implementations against a range of identified criteria. The full study, "In-nimbo sandboxing," can be accessed at http://doi.acm.org/10.1145/2600176.2600177. For the SoS newsletter article, visit http://www.cps-vo.org/node/21422.

[Photo credits: Lordn, mipan/iStock/Thinkstock]

Science of SecUrity and REsilience for Cyber-Physical Systems project

On 17–18 March 2015, the NSA Trusted Systems Research Group met in Nashville, Tennessee with academic researchers from the Science of SecUrity and REsilience for Cyber-Physical Systems (SURE) project to review their first six months of work. Researchers came from four participating institutions—Vanderbilt University, the University of Hawai'i, the University of California at Berkeley, and the Massachusetts Institute of Technology (MIT). Individuals from the National Science Foundation, Nuclear Regulatory Commission, and Air Force Research Labs also attended.

SURE is the NSA-funded project aimed at improving scientific understanding of "resiliency"; that is, the robustness of a cyber-physical system (CPS) to reliability failures or faults, as well as survivability against security failures and attacks. Initially, SURE focused on CPS architectures related to only water distribution and surface traffic control; the project now also focuses on air traffic control and satellite systems. The principal investigator for the SURE project is Xenofon Koutsoukos, professor of electrical engineering, computer science, and senior research scientist in the Institute for Software Integrated Systems at Vanderbilt University. Professor Koutsoukos indicated the use of these additional CPSs is to demonstrate how the SURE methodologies can apply to multiple systems.

The SURE project addresses the question of how to design systems that are resilient despite significant decentralization of resources and decision-making. Main research thrusts include hierarchical coordination and control, science of decentralized security, reliable and practical reasoning about secure computation and communication, evaluation and experimentation, and education and outreach.

The Resilient Cyber Physical Systems (RCPS) test bed, discussed in Emfinger, Kumar, and Karsai's article in this issue ("Resilient and secure cyber-physical systems"), supports evaluation and experimentation across the entire SURE research portfolio.

In addition to the RCPS test bed, researchers presented 10 resiliency projects on behavioral and technical subjects including adversarial risk, active learning for malware detection, privacy modeling, actor networks, flow networks, control systems, software and software architecture, and information flow policies. The scope and format of the Cyber Physical Systems Virtual Organization web site (http://cps-vo.org) was also briefed. For the complete Science of Security newsletter article, visit http://www.cps-vo.org/node/21425.

[Photo credit: Nerthuz/iStock/Thinkstock]

Selecting Android graphic pattern passwords

With funding from the University of Maryland Science of Security Lablet, researchers at the US Naval Academy and Swarthmore College conducted a large study of user preferences relating to usability and security of graphical password patterns used to access Android mobile phones.

Android mobile phones come with four embedded access methods: a finger swipe, a pattern, a PIN, or a password, in ascending order of security. In the pattern method, the user is required to select a pattern by traversing a grid of 3 x 3 points. A pattern must contain at least four points, cannot use a point twice, and all points along a path must be sequential and connected; that is, no points along the path may be skipped. The pattern can be visible or cloaked.

When a user enables such a feature, however, how does that person trade off security with usability? Also, are there visual cues that lead users to select one password over another and whether for usability or security?

The study by Adam Aviv and Dane Fichter, "Understanding visual perceptions of usability and security of Android's graphical password pattern," uses a survey methodology that asks participants to select between pairs of patterns and indicate either a security or usability preference. By carefully selecting password pairs to isolate a visual feature, a visual perception of usability and security of different features were measured. The 384 users in the study sample were self-selected via Amazon Mechanical Turk, an online marketplace for crowdsourcing. Users found that visual features that can be attributed to complexity indicated a stronger perception of security, while spatial features, such as shifts up/down or left/right are not strong indicators either for security or usability.

In their study, Aviv and Fichter selected pairs of patterns based on six visual features:

1. **Length:** Total number of contact points used
2. **Crosses:** The pattern double-backs on itself by tracing over a previously contacted point
3. **Nonadjacent:** The total number of nonadjacent swipes which occur when the pattern double-backs on itself by tracing over a previously contacted point
4. **Knight-moves:** Moving two spaces in one direction and then one space over in another direction, like the knight piece in chess
5. **Height:** The amount the pattern is shifted towards the upper or lower contact points
6. **Side:** The amount the pattern is shifted towards the left or right contact points

Users were asked to choose between two passwords, indicating a preference for one password meeting a particular criterion (such as perceived security or usability) over the other password. By carefully selecting these password pairs, researchers could isolate the passwords' visual features and measure the impact of a particular feature on users' perception of security and usability.

The researchers concluded that spatial features have little impact. More visually striking features have a stronger impact, with the length of the pattern being the strongest indicator of preference. These results were extended and applied by constructing a predictive model with a broader set of features from related work, and they revealed that the distance factor, the total length of all the lines in a pattern, is the strongest predictor of preference. These findings provide insight into users' visual calculus when assessing a password, and this information may be used to develop new password systems or user selection tools, like password meters.

Moreover, Aviv and Fichter conclude that, with a good

[Photo credit: Kirillm/iStock/Thinkstock]

predictive model of user preference, this research could be expanded and these findings could be applied to a broader set of passwords. For example, ranking data based on learned pairwise preferences is an active research area in machine learning, and the resulting rankings metric over all potential patterns in the space would greatly benefit the security community. For example, such a metric could enable new password selection procedures. The full study is available at http://www.usna.edu/Users/cs/aviv/papers/p286-aviv.pdf. For the complete Science of Security newsletter article, visit http://www.cps-vo.org/node/21423.

Computational Cybersecurity in Compromised Environments Workshop

The Special Cyber Operations Research and Engineering (SCORE) subcommittee sponsored the 2014 Computational Cybersecurity in Compromised Environments (C3E) Workshop at the Georgia Tech Research Institute (GTRI) Conference Center in Atlanta from 20–22 October 2014. This event, the sixth in a series of annual C3E workshops, brought together top academic, industry, and government experts to examine new ways of approaching the nation's cybersecurity challenges. In particular, participants discussed how to enable smart, real-time decision-making in situations involving both "normal" complexity and persistent adversarial behavior.

Since its start in 2009, C3E's overarching objectives have been to develop ideas worth additional research and to develop a community of interest around unique, analytic, and operational approaches to the persistent threat.

C3E 2014 continued to focus on the needs of the practitioner and leverage past themes of predictive analytics, decision-making, consequences, and visualization. The two tracks, Security by Default and Data Integrity, were developed based on recommendations from senior government officials in order to ensure that outcomes will be applicable to real-world security needs. As in previous years, the event included keynote addresses on a wide range of cybersecurity topics, a discovery or challenge problem that attracted participants' analytic attention prior to the workshop, and track work that involved small-group focus on security by default and data integrity.

This year's discovery problem focused on approaches that challenge traditional thinking on using metadata to identify high-interest, suspicious, and likely malicious behaviors. Participants used the Department of Homeland Security's Protected Repository for the Defense of Infrastructure against Cyber Threats (PREDICT), a rich collection including the borderline gateway protocol, the domain name system, data applications (net analyzer), infrastructure (census probe), and security (botnet sinkhole data, dark data, etc.).

The Data Integrity Track addressed issues associated with finance and health/science, and captured relevant characteristics. Track participants also identified potential solutions and research themes for addressing data integrity issues: 1) diverse workflows and sensor paths, 2) cyber insurance and regulations, and 3) human-in-the-loop data integrity detection.

The Security by Default Track focused on addressing whether it is possible to create systems that are secure when they are fielded, and examined the common perception among stakeholders that such systems may be less functional, less flexible, and more difficult to use. Participants identified five topics that might merit additional study:

1. The "building code" analogy—balancing the equities,
2. Architecture and design—expressing security problems understandably,
3. Architecture and design—delivering and assuring secure components and infrastructure,
4. Biological analogs—supporting adaptiveness and architectural dynamism, and
5. Usability and metrics—rethinking the trade-off between security and usability.

For the complete article, visit http://www.cps-vo.org/node/21421.