The Resaerch Value of Publishing Attacks

David Basin and Srdjan Capkun
Institute of Information Security
ETH Zurich

December 20, 2012

Information Security is booming. Companies are making money selling fear and countermeasures. The research community is also extremely active, churning out papers featuring attacks on systems and their components. This includes attacks on traditional IT systems as well as on IT-enhanced systems, such as cars, implantable medical devices, voting systems, and smart meters, which are not primarily IT systems but have increasing amounts of IT inside. Moreover, any new paper on analysis methods for critical systems is now considered incomplete without a collection of security-relevant scalps on its belt. Pretty much every system imaginable, critical or not, is now a target of attacks.

There are good reasons for this trend. Fear sells! Headlines are good for conference attendance, readership, and tenure cases. Moreover, negative messages about successful attacks are simple and understandable by the general public, much more so than other research results. And security and insecurity are, after all, two sides of the same coin.

Seek and ye shall find

Systems have bugs and large, complex systems have many bugs. In their recent analysis of open source projects, [Cov11] used a static analysis tool to find 16,884 defects in ca. 37.5 million lines of source code from well-managed open source projects, which is approximately 0.45 bugs per 1000 lines of code. These were medium to high risk defects, including typical security-critical vulnerabilities such as memory corruption problems and API usage errors. For large-scale projects, developers cope with the seemingly infinite number of bugs in their products by employing triage processes to classify which bugs they work on and which they ignore. There are simply too many to address them all.

This should not come as a surprise. Complexity is at odds with security. Moreover, economic factors are often at play, where timeliness and functionality are more important than security. But there are other reasons too why insecurity is omnipresent.

To begin with, systems undergo constant evolution. There has been a recent surge in attacks where once-closed systems, like medical devices and cars, open

up and are enhanced with new communication interfaces (see e.g., [HHBR+08, RMM+10, FDC11]). The problem here is that the extended capabilities were usually not anticipated in the original design, often resulting in vulnerabilities that are easy to exploit. Not surprisingly, adding wireless communication without measures to ensure the confidentiality and authenticity of transmitted data results in a system vulnerable to eavesdropping and spoofing. This problem is particularly acute for products manufactured by traditional industries that did not previously require expertise in Information Security.

Systems not only interface with the outside world, they also interface with each other. For their composition to be secure, the assumptions of one subsystem must match the guarantees of the other. However, economics and market availability often dictate the choices made, especially for hardware components where manufacturing one's own components is often not an option.

Finally, even when a system's security is carefully analyzed, this analysis depends on the deployment scenarios considered, in particular, the associated adversary model. What kind of an adversary should the system withstand? A system secure against a network attacker may be completely insecure against one with a screw driver and physical access to the server. Many IT-enhanced systems have been developed using proprietary protocols and communication technology, leading to the belief that it was difficult for outsiders to interface with them. However, for wireless communication, the increasingly wide-spread availability of tools and equipment, such as Universal Software Radio Platforms, has made it easy and inexpensive for nonspecialists to communicate with even the most exotic systems, thus dramatically changing the adversary's capabilities. As scenarios and adversaries change over time, so do the possible attacks.

Summing up, it is not surprising to see so many system attacks reported, in particular on IT-enhanced systems. But what makes attacks worthy of scientific publication? Are all these attacks of the 'yet another buffer overflow' variety? Is there any point in publishing research papers that feature attacks on systems that were not designed to resist attacks, not used as they were designed, or used in scenarios for which they were not designed?

Learning from attacks

A hallmark of good research is the generality of the insights gained. In security, these are insights into the problem and countermeasures.

Increasing awareness is a common argument for publishing attack papers and has its merits. In particular, a heightened awareness of problems and their severity may lead to the system in question being withdrawn from service; alternatively, others can follow up with designs that solve the documented problems. Such attacks have, in the past, raised awareness among policy makers of the immaturity of existing technologies and the associated risks. This is particularly valuable for new systems and technologies. Here, the novelty of the kind of attack is less relevant than the novelty of the system and the impact of its compromise.

Although raising awareness is important, it can backfire as too much sensationalism numbs the readers' sensitivity to what the real problems are. And there is usually limited research value in just showing how standard problems can be exploited in yet another setting. It is clear that unauthenticated communication opens the door to spoofing attacks, whether we are talking about cars, medical implants, or personal robots. The same holds for standard, well-studied, software vulnerabilities. In contrast, a paper that refines an existing attack, demonstrates a novel kind of attack, or contributes to new attacker models can have tremendous research value.

One benefit of studying attacks is a better understanding of the cause of the underlying vulnerability, for example, whether it is the result of a design or implementation error, the unavailability of solutions on the market, improper system usage, or an oversight in the risk analysis. This last reason occurs surprisingly frequently; systems are often left unprotected because the designers simply do not believe that they need to be protected or assume that the systems are sufficiently closed or obscure and therefore unlikely to be reverse-engineered by attackers (or determined researchers). As recent attacks on medical devices and modern cars show, these assumptions are incorrect.

An attack paper can also explicate what is required for a successful attack. Is the exploitation of a vulnerability straightforward or only possible by well-funded, technically sophisticated attackers? The devil is in the details! A good attack paper can show how to construct an exploit and the cost of doing so. Moreover, it can help refine the conditions under which the attack may succeed and its success probability. For example, an attack might be conditioned on the attacker's physical location, antenna size, transmission power, etc. For example, the success of spoofing attacks on Global Positioning System receivers strongly depends on the locations and characteristics of the attackers' antennas.

Let us expand upon this last point. What makes security special is the role of the adversary. A system's security can only be evaluated with respect to a model of the adversary, i.e., a description of his capabilities. Thus, in our view, the most important reason for studying attacks is that they can help refine this model for the domain at hand. Below we give two examples of this from the domain of security protocols and relay attacks.

In 1978, Needham and Schroeder proposed one of the first authentication protocols. Their protocol used public key cryptography to achieve mutual authentication between two principals in the presence of an attacker who can eavesdrop and spoof messages. 18 years after its publication, Lowe [Low96] showed that the protocol could be attacked by a man-in-the-middle, who executes the protocol as an insider in two interleaved sessions. This attack sensitized the security protocol community to the importance of considering adversaries who have insider capabilities. Later, motivated by attacks on long-term keys stored in memory, weak random number generators, and the ability of adversaries to read out part of an agent's session state, cryptographers developed a host of more refined adversarial models and security definitions reflecting these increased capabilities. These new models have led to improved protocols as well as methods and tools for reasoning about the security of protocols and systems,

with respect to these refined adversarial models, see e.g. [BC10].

A second, more recent, example are Relay, Mafia-Fraud and Wormhole attacks where the attackers simply relay messages, unmodified, between the two communicating parties. Such attacks have been recently used to attack entry and start systems in cars [FDC11] and payment systems that rely on near field communication. These attacks showed that the success of relay attacks on such systems strongly depends on the speed that attackers can process signals. They further demonstrated that existing technology enables attackers to build relays that have practically undetectable processing delays. This was particularly important in the case of entry and start systems for cars; the attacks revealed that these systems can only detect relays that introduce delays longer than several microseconds. This led to refined attacker models and also motivated new security solutions, for example distance bounding protocols.

Final thoughts

As our physical and digital worlds become more tightly coupled, the incidence of attacks will increase as well as their consequences. Many of these attacks will be news-worthy, but most will not be research-worthy. This does not mean that papers featuring attacks on highly visible systems should not find their way into research conferences; having had such papers published, the authors of this column do appreciate that the community accepts results of this kind. However, as researchers we should have high aspirations. With every attack paper there is an opportunity to truly contribute to the community with new insights into both systems and their vulnerabilities, and adversaries and their capabilities. We believe that one should take this opportunity and, after discovering an attack, take a step back and reflect on what can be learned from it, and afterwards present it to the community.

References

[BC10] David Basin and Cas Cremers. Modeling and analyzing security in the presence of compromising adversaries. In *Computer Security - ESORICS 2010*, volume 6345 of *Lecture Notes in Computer Science*, pages 340–356. Springer, 2010.

[Cov11] Coverity scan: Open source integrity report, 2011.

[FDC11] Aurélien Francillon, Boris Danev, and Srdjan Capkun. Relay attacks on passive keyless entry and start systems in modern cars. In *Proceedings of the Network and Distributed System Security Symposium*, 2011.

[HHBR+08] Daniel Halperin, Thomas S. Heydt-Benjamin, Benjamin Ransford, Shane S. Clark, Benessa Defend, Will Morgan, Kevin Fu,

Tadayoshi Kohno, and William H. Maisel. Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses. In *Proceedings of the 2008 IEEE Symposium on Security and Privacy*, SP '08, pages 129–142, Washington, DC, USA, 2008. IEEE Computer Society.

[Low96] Gavin Lowe. Breaking and fixing the Needham-Schroeder public-key protocol using FDR. *Software — Concepts and Tools*, 17(3):93–102, 1996.

[RMM+10] Ishtiaq Rouf, Rob Miller, Hossen Mustafa, Travis Taylor, Sangho Oh, Wenyuan Xu, Marco Gruteser, Wade Trappe, and Ivan Seskar. Security and privacy vulnerabilities of in-car wireless networks: a tire pressure monitoring system case study. In *Proceedings of the 19th USENIX conference on Security*, USENIX Security'10, pages 21–21, Berkeley, CA, USA, 2010. USENIX Association.