Abstract: Regulatory Options for the Global Trade in Zero-Day Vulnerabilities

Mailyn Fidler | Stanford University | mfidler@stanford.edu

The emergent "gray-market" global trade in zero-day vulnerabilities involves governments purchasing vulnerabilities for cyber attack and espionage purposes. This market poses negative security consequences and raises questions about what regulatory mechanisms might best address these consequences.

Zero-day vulnerabilities are flaws discovered in existing software programs about which neither the responsible company nor the public knows. These vulnerabilities are exploitable on the "zero-th" day of their existence, and their secrecy and immediate exploitability make them valuable tools. Zero-days were previously the domains of software security researchers, who reported zero-days to responsible companies for free or a small reward, or of hackers, who profited by selling them on the black market for largely criminal purposes.

Governments, particularly the U.S. government, have started purchasing zero-day vulnerabilities for use in cyber attack and cyber espionage, paying high prices, building a stockpile, and feeding a thriving gray market. Prices on the gray market range from about $16,000 to $250,000 per vulnerability, usually much higher than prices on the black and white markets. The NSA employs in-house researchers to find zero-days, but the U.S. government also allocated $25 million for purchase of zero-day vulnerabilities in fiscal year 2013. The U.S. government used zero-day vulnerabilities in Stuxnet, the cyber attack against Iranian centrifuges, and in programs such as NSA's FoxAcid, which compromises targeted computers.

Although some contend the U.S. government has compelling national security reasons to participate in the zero-day market, others criticize the practice for its broader cybersecurity consequences. The current public understanding of U.S. government policy is that the government does not notify affected companies about vulnerabilities it identifies or purchases. This practice leaves companies and citizens vulnerable to exploitation if other parties discover the flaw, which undermines citizen cybersecurity in pursuit of other national security objectives. The success of government identification, purchase, and deployment of zero-day vulnerabilities depends on the continued vulnerability of everyone else. Similarly, high gray market prices divert trade from the white market, making the white market less lucrative than when it only competed with the black market. On an international level, the burgeoning gray market means U.S. adversaries with low cyber capacities can access "ready made" cyber attack tools, potentially more rapidly achieving the capability to threaten U.S. interests in cyberspace.

Given these negative consequences, my research investigates options for regulating the zero-day gray market. Examining both domestic and international approaches, I analyze a suite of tools ranging from "soft" to "hard" law. On the domestic side, I examine criminalization, U.S.-based export controls, and inter-agency transparency-building initiatives. On the international side, I analyze potential initiatives within existing international organizations, non-binding but normative restrictions on exports through the Wassenaar Agreement, and the possibility of a binding treaty. Currently, my research demonstrates that each option has significant drawbacks, but these options are part of ongoing policy discussions. Analyzing the potential and downsides of each option is intended to serve as a useful resource for policymakers.