The limits of measurement for cyber security

A fundamental limitation to all scientific research is the requirement for informed consent. In the computer security arena, this is even more problematic because (1) we often using the techniques of psychology requiring deception to get valid results, (2) we are measuring things that, in many cases, the people involved might/would not give consent to measure, (3) there is a lack of repeatability, (4) when we measure we may disturb the environment to the point where the measurements are no longer suitable to their purpose, and (5) permission may itself invalidate the results by, among other things, producing self-selected sample sets.

While technical measurement of things like known attacks and social behaviors may avoid many of these issues, and broad sweeping consent associated with click contracts may absolve researchers of some technical legal issues, issues associated with influence operations, deception and counter-deception, artificial amplification, and controls over meme spread are often problematic. From a legal perspective, researchers are not somehow immune from the law, whether criminal or civil, because they are government funded. Mere measurement of spread is inadequate to measure the effect of countermeasures. Observing traffic ignoring content is problematic for analysis of memes. Countermeasures may be associated with interference, and certainly have effects on the environment they are experimenting on. Repeatability issues abound, User interfaces and many supporting systems are often fragile and the systems they are part of may be brittle, leading to induced failures by experimentation. Measurement by induced and suppressed signal techniques may cause cognitive effects on individuals and groups that last far beyond the period of the experiment. The list goes on and on.

Another critical issue is the difference between what is a meaningful measure and what we may be able to measure. In many cases, researchers have sought to measure one thing or another, but done so by using surrogates that may not reflect the phenomena of interest. Even the basic of measurement theory are often ignored in such research. If we are measuring human behavior on large scale, how do we do so including the non-online behaviors. In the cyber-security arena, attackers identify defenses and detection mechanisms and seek to avoid them. Thus the things we measure may not indicate the changes in attacker behavior associated with their detection of the measurement mechanisms. Many network-related effects involve behaviors not contained with in the networks themselves. The spread of memes may and often do leave the networked environment on one place and re-enter elsewhere.

Even the most basic measurements in common use, such as prevalence of detected vulnerabilities, quantities of known attack-associated events, and indicators of compromise, end up applied without clarity around the consequences of those attacks if successful against their targets. And essentially all such methods ignore the increasing use of deception for defense, including false positive generations and adaptive defenses that alter the behavior of systems in response to detections. The measurement mechanisms themselves may also produce vulnerabilities in systems, particularly in the case of active measurements which may be exploited in a variety of ways. Thus there are questions of “do no harm” associated with measurement methodologies.

Finally, there is a barrier to gaining access to real data because of a wide range of concerns about privacy, the extensive use of encryption, the use of load balancing and address translation, high bandwidth usage, low signal to noise ratio, and related issues with sensor placement and distribution of analytical capabilities.

The question version of these challenges are, in essence, how do we mitigate these limitations and create measurement infrastructures and methodologies that allow researchers to systematically do so without spending inordinate amounts of time and effort? How do we standardize such an infrastructure? How do we teach researchers about these issues and how to address them? And how do we do so without substantial added expenses? Or in fact do we take the approach of most other areas of science and educate researchers in these issues and embed them in the basics of how to do research in the cyber era?