

GRANDON GILL

EXPANDING JOINT VULNERABILITY ASSESSMENT BRANCH¹

If we needed to hire a new team member or two, we could probably handle that. But if you ask us to double in size, it would take years. What can we do if the demand for our services keeps expanding?

David Rohret, the founder of the Joint Vulnerability Assessment Branch (JVAB) pondered this difficult question. Since 2003, he had been involved in building a team that was uniquely positioned to identify a wide range of vulnerabilities in military and commercial communications and web-based systems. He could cite numerous examples of past situations where the early use of JVAB's services led to, or could have led to, tens of millions of dollars in savings—or possibly more, had the issues they detected been left unattended. The value that JVAB offered was gradually being recognized and, as a result, demand for their services was building. The problem was that it was nearly impossible to hire people with the skills necessary to meet the growing need.

There were a number of aspects of JVAB's approach that made it unique. First and foremost, it had been early to recognize that formerly distinct elements of communications systems were rapidly converging. Historically, communications using radio frequency (RF) signals had been the domain of electrical engineers, while network communications were handled by computer scientists. As network traffic was increasingly being handled using cellular and wifi signals, however, RF intrusions became a serious threat. By the same token, RF communications—such as those handled using high end hand-held devices and cell phones—often relied on the same IP protocols used by the Internet—making them a potential pathway to servers.

Another key aspect of JVAB was its adversarial mindset. It prided itself on using the same tools and techniques as the black hat hackers that threatened systems in real world settings. Not only was this an attitude that was generally not cultivated in educational institutions, it also ran counter to the experience of individuals that has spent all their professional life dealing with security in a defensive posture.

In the past, Rohret had hired high potential individuals, usually with military experience, and had helped them develop their skills over many years. The end result was the formation of a team with an extraordinary track record of success. But if JVAB were to meet the continuing demand for its services, it needed to figure out new ways to expand. That could be a real challenge in an organization where the most valuable assets all wore shoes.

¹ Copyright © 2016, *Grandon Gill*. This case has been reprinted from the *Muma Case Review*, Volume 1, Number 5 and was prepared for the purpose of class discussion, and not to illustrate the effective or ineffective handling of an administrative situation. Names and some information have been disguised. This case is published under a Creative Commons BY-NC license. Permission is granted to copy and distribute this case for non-commercial purposes, in both printed and electronic formats.

Introduction to Hacking

In the early days of computing, the term hacking referred to the ability to code, navigate and manipulate information and workaround problems presented by complex information and communications systems. A “hacker” referred to an individual skilled in the art of hacking. Over time, however, the term began to acquire a more pejorative meaning in common usage and was more often applied to individuals who broke into systems with malicious intent—also sometimes referred to as “crackers”.

Hacking Skills and Personas

Negative connotations aside, effective hacking necessarily involved an admirable set of skills. Coding skills were required to identify potential vulnerabilities in systems that could be exploited for the purpose of intrusion. This knowledge must include a deep understanding of machine architecture and data representation. Familiarity with different operating systems was required to navigate within and across systems. Expertise in applications software, particularly database systems, was critical, as was experience with alternative network architectures, communications protocols and alternative modes of data transmissions. And, of course, familiarity with a wide range of encryption techniques and security procedures was a must.

In the early days of computing, it was occasionally possible for an individual to acquire a complete portfolio of skills such as those described above. With the growth of the Internet and the explosion in the variety of technologies in broad use, no single individual was likely to possess all the skills necessary to be an all-in-one hacker. As a consequence, successful hackers tended to share skills and tools. This could be accomplished by working in teams consisting of different specialists or by sharing toolboxes on black hat forums and chat rooms. For example, a hacker seeking to deploy malware on an unsuspecting organization’s network would not necessarily have to write the malware himself or herself. Instead, he or she might communicate with other hackers online looking for an existing “packaged” solution. In that way, the acquisition process would not be so different from that of an organization’s IT group seeking to find a software solution to a business problem.

There was, however, an important distinction that must be made between commercial solutions and hacking solutions. In most cases, neither the buyers nor the sellers of commercial software particularly cared if the transaction became known by others. In the hacker community, however, the situation was very different. Much of the technology and information exchanged by hackers was either itself illegal, or was most suited for illegal purposes. Naturally, neither the supplier nor the acquirer wanted anyone to know who participated in the transaction. For this reason, a hacker would typically cultivate one or more personas—screen names that disguised the individual’s true identity.

Building a persona was a time consuming pursuit. Because the importance of hacker sites had not gone unnoticed by law enforcement, the natural reaction of experienced hackers was to be extremely suspicious of new entrants to a hacker site. For this reason, it could take several years before a particular individual was sufficiently trusted so that any information or tools that were not widely known could be acquired from members of the site. Moreover, because anonymity was prized by the participants on such sites, it was implicitly understood that any profile information supplied by members was likely to be misleading, if not wholly false. Furthermore, should an individual’s identity be compromised—e.g., be discovered to be a member of law enforcement or employed by a U.S. government agency—he or she was likely to be “outed” by other members. As a consequence, individuals seeking to penetrate the hacker community would often establish multiple personas simultaneously. Doing so allowed them to continue participating in the site even after one of more of their identities had been compromised.

White Hat vs. Black Hat Hackers

The superior hacker must possess not only formidable technical skills but also personality traits that included creativity, curiosity and persistence. It was not surprising, therefore, that the same skill set was highly prized by organizations whose goals were not necessarily nefarious—such as companies and the government. Because the term hackers has acquired a negative connotation, a distinction became commonly made between “white hat” and “black hat” hackers.

Black hat hackers tended to live up to the negative stereotypes associated with hacking, employing their skills in criminal acts, either for personal gain or simply to create mayhem. In direct contrast, white hat hackers employed their skills in ethical computing, and generally followed a code such as the *Association for Computing Machinery’s* 1992 Code of Ethics and Professional Conduct (see Exhibit 1 for the ACM Code’s section on moral imperatives). These individuals often had jobs in cybersecurity, testing and systems design, where their skills were critical for ensuring that systems were not unnecessarily vulnerable to black hat hacking.

Between black hat and white hat hacking, there was a middle ground that was, unsurprisingly, sometimes referred to as “grey hat” hacking. These individuals may sometimes employ hacking for personal gain and, occasionally, may choose to skirt the law. They may also switch from black hat to white hat hacking activities over time. A number of well-known hackers graduated from black hat hacking to fighting black hat hackers. Kevin Mitnick, for example, was forced to spend several years in prison for his computer crimes but, upon his release, he went on to launch a well-respected computer security firm. Indeed, the process of establishing a credible persona within the black hat hacker community may require the white hat hacker to participate in some activities that were more than a little bit grey. How these actions were classified would likely depend on the hacker’s intent. To clarify this, it was useful to introduce the concept of red and blue teams.

Red Team and Blue Team

The skill set of white hat hackers could serve two particularly valuable purposes within an organization:

1. To ensure the design of systems was as well defended against intrusion and compromise as possible, and
2. To test the defenses provided by those same systems in sophisticated and creative ways, a process referred to as penetration testing

The first of these activities was primarily defensive. The second involved coming up with attacks; essentially playing offense.

In military exercises, the defensive group was typically referred to as the *blue team*. The offensive group—which effectively attempted to mimic the behavior of black hat hackers—was referred to as the *red team*. There was considerable overlap in the skills required for the two teams. They normally involved separate groups, however. Having designed a system’s defenses, the blue team was often unaware of potential weaknesses. (Indeed, had they been aware of security holes, they would have patched them). Thus, the most effective attackers were those without preconceptions of the measures in place to prevent their penetration.

Rohret and the JVAB team normally assumed a red team role in their activities. In doing so, they had developed a framework that defined the Adaptive Red Team (ART). The key elements of this framework methodology were summarized in Exhibit 2. Intentionally broad, the framework included operational,

technological and environmental considerations. It also did not limit itself to uncovering vulnerabilities. Specifically, the methodology stated:

In addition to identifying exploitable IRC vulnerabilities, the ART will research potential mitigation efforts that counter identified vulnerabilities, using both passive and active measures.

In other words, under the ART framework, the red team would also undertake some activities more traditionally associated with the blue team.

Joint Vulnerability Assessment Branch

The Joint Vulnerability Assessment Branch (JVAB) was founded in 2003 under the auspices of the Office of the U.S. Secretary of Defense (OSD). Its mission was to provide an agile assessment of the security of a wide range of technologies that were deployed or were being considered for deployment in the field. In the past, such assessment had required years to conduct. At that time, the world was less connected and many of the technologies involved were specifically created to military specifications. Unfortunately, as more and more off-the-shelf technology solutions were being adopted by military and disaster recovery agencies, such delays were no longer acceptable. If security was to be maintained, time horizons of months, not years, had to be achieved.

In establishing the group, its founder David Rohret, sought to break down established approaches to addressing security:

In the past, cybersecurity was handled by three distinct groups: operators, computer scientists and RF [radio frequency] engineers. The operators looked at system security from the user's perspective, focusing on the interface. The computer scientist looked at the system from the perspective of logical design and code security. The RF engineers looked at the transmission of signals, with a particular emphasis on interference.

This separation did not serve us well at the time of the case. For example, the old belief was that RF was little more than energy. In current environments, however, RF carried IP [internet protocol] packets, making it vulnerable to types of intrusion and spoofing that could not even be imagined in old voice-only systems. Think about how different smart phones were at the time of the case from the old analog land lines.

What was needed, therefore, was a team encompassing the set of all three skills. That team needed to adopt an adversarial point of view. It had to focus on the *goal* of penetrating the system being tested, no matter how that goal was best achieved, rather than focusing on the vulnerabilities of the individual *technologies* used to construct the system.

Examples of JVAB Activities

The efficacy of JVAB's approach to red team activities had been demonstrated on many occasions. Some of the many examples Rohret provided included:

- Breaking through the online security of a well-known defense laboratory in under 20 minutes.
- Disrupting communications with a precision guided surveillance system prototype that had cost \$20 million to build. The organization and a well-known university that had developed the prototype had refused an offered pre-assessment. What JVAB had discovered, however, was that the wireless communication protocol being used was so sensitive that taking out a communications

packet by inserting just one bit every five or six seconds was sufficient to render the system inoperable.

- Breaking into a secure facility using a semi-autonomous vehicle. While the facility had sensors and automated weaponry worthy of a Hollywood action film, JVAB was able to identify a key IP address through which it was able to disable the gate, spoof the sensors and, effectively, neutralize the guards.
- Determined a vulnerability through which a Southeast Asian terrorist group was taking down U.S. satellite reconnaissance. What JVAB discovered was that the satellites had a number of unused communication channels, left open for future use. Through these channels, it found Cisco routers that had been left in debugging mode by default—effectively disabling the routers’ built-in security. Through these routers, JVAB was able to assume control of the satellite.

These examples, and many others, highlighted the value of the multidisciplinary goal-focused strategy employed by JVAB. He perceived the need for this type of approach to be ever-growing. As an example, he stated:

90% of that technology, and of what we [JVAB] do, is unclassified. Even hobbyist sites contain all sorts of software for doing things... Governments are often at a disadvantage here. Thailand does not have the money to acquire the technology that its cocaine dealers buy.

JVAB Operations

JVAB was a small organization, typically running with 6-9 employees, nearly all of whom had substantial experience in one or more of the three cybersecurity areas. Although it was headquartered near San Antonio, Texas, most of their activities involved travel to remote sites. For this reason, Rohret estimated that the group spent over 50% of their time on the road.

An example of JVAB’s operations was the services that it provided to the *Joint Interagency Field Experiments (JIFX)* event hosted by the *U.S. Naval Postgraduate School* (Murphy, et al., 2015, p. 63), described as follows:

The Joint Interagency Field Experimentation (JIFX) event, organized by the Naval Postgraduate School (NPS), is conducted 3-4 times a year at various locations. The four-day event [is] specifically designed to facilitate structured and unstructured communications between a variety of parties—e.g., software developers, inventors, military and civilian users of various technologies, academics, and agencies responsible for identifying and procuring technology solutions—that frequently are constrained in their informing activities in more restrictive venues. Over the course of the event, participants may observe technology demonstrations, obtain feedback from potential users, acquire new ideas about their technologies might be employed and, perhaps most significantly, engage in ad hoc collaborations with other participants.

Many of the technologies that were demonstrated and employed in the event’s many experiments involved information and/or communications technologies. JVAB was retained by the event to test these technologies at no cost to participants. It performed two distinct types of tests:

1. *RF testing*. Done to ensure that any technologies that radiated electromagnetic waves did so at the proper frequencies, so as to avoid interference. This was a critical concern at JIFX because many of the experiments involved unmanned aerial vehicles (UAVs) and interference could lead to

crashes. Improper broadcasts could also interfere with the operations on the bases where JIFX was held.

2. *Vulnerability assessment.* Consistent with JVAB's principal mission, Rohret and his team performed a broad array of penetration tests and disruption tests on both information and communications systems. Using a large set tools acquired on hacker sites and some very expensive communications equipment, the team could run programs, websites, and many types of equipment through a series of tests. Rohret estimated that the same level of testing—particularly for complex systems that were in operation or nearly in operation—would cost tens (or even hundreds) of thousands of dollars if purchased from a commercial provider.

Because of the high value provided by JVAB's security assessments, Rohret received frequent requests for JVAB's services. This created a serious challenge for the organization. As previously noted, the team already spent most its time on the road. Given the personnel-intensive nature of its services, that precluded taking on many additional engagements. If JVAB were to expand, it would therefore need to add a significant number of individuals to its ranks, permitting it to field multiple teams. The question then became: how do we acquire these individuals?

Challenge of Recruiting for JVAB

Even before JVAB's special needs were considered, acquiring capable cybersecurity professionals was rapidly becoming a crisis for government and industry. A 2015 Frost and Sullivan study of the information security workforce began with the following statement (Suby & Dickson, 2015, p. 3):

The information security workforce shortfall is widening. In this year's survey, 62% of the survey respondents stated that their organizations have too few information security professionals. This compares to 56% in the 2013 survey. Also in a shift from the 2013 survey, the reasons for this hiring shortfall are less about money as more organizations are making the budgets available to hire more personnel. Rather, an insufficient pool of suitable candidates is causing this shortfall.

Beyond this general shortfall, however, the unique nature and strategy of JVAB made finding suitable candidates even more problematic. Rohret described a number of the issues JVAB had faced:

One problem is finding individuals with RF expertise. There used to be a ton of RF experts, back when most communications were analog. Today, with the growing emphasis on digital communications, there are very few. Those that are available command very high salaries--\$160,000 a year is not unusual. Unfortunately, we need this expertise. As I mentioned before, digital communication moves over RF carriers; that can be a major source of vulnerability that we cannot afford to ignore...

Computer science experts are hard enough to find. But we need even more. They need to have an adversarial mentality to be effective as members of a red team. The problem is that adversarial hackers often have a police record. The nature of our government clientele means that we cannot hire these individuals, even if they have reformed...

Hiring individuals directly out of school also presents obstacles. For one thing, even the best cybersecurity programs—and there are not many of these—generally teach material that is at least five years behind the times. Even if we were to train these individuals, we need them to have experience in defense before we update their attack skills. This needs to be combined with a computer science or electrical engineering degree, unless the candidate has very specialized experience (such as might be acquired in the military)... We have also found that certifications are not

all that useful. What they cover is too simple for our purposes and is also likely to be out-of-date...

There is also the issue of having an uncompromised persona. Without a valid persona you cannot gain access to the newest and most effective tools in the black hat hacker's arsenal. These take years to develop, however, and a single slip can render them useless. Over my career, I have developed ten or more of these. Most were eventually compromised... We had one applicant come in who assured us that he had a valid persona. I recognized the name and showed him a site where it had been "outed". And that had happened years before he came to us...

The hiring situation was not altogether bleak. In their favor, Rohret and his team members had an immense network of contacts. JVAB's reputation and this network ensured that some qualified applicants would always find them. But reputation was a double-edged sword. Too often, members of the team would find themselves being approached by other firms—often offering twice what JVAB could afford to pay. As a result, a certain amount attrition was inevitable and some recruiting was needed just to maintain the status quo.

Potential Alternatives

Rohret did not feel any compelling need to turn JVAB into a giant operation. Nevertheless, he felt that the assessments that they provided were a major service to private organizations and the government. He wanted to ensure that they continued. In thinking about possible directions, he looked at his options from both short and long term perspectives.

Short Term "Band Aids"

As a consequence of JVAB's network of connections, Rohret felt that making one or two hires in the short term was a reasonable expectation. There were also a few things that they might do to leverage existing personnel and, perhaps, to modify their business model. He viewed these solutions as short term "band aids", some of which would be necessary if operations were to be maintained and, perhaps, somewhat expanded.

Assigning Personnel by Technology Readiness Level (TRL)

One possible approach to leveraging additional personnel being considered by JVAB would take into account technology readiness level (TRL). The TRL scale was based upon how close a particular technology or system was to active use. Low TRL systems were those in the early prototype stage. High TRL systems were at late stage testing (e.g., beta tests) or were actually in use. The security needs of the various TRL levels could differ markedly.

At the low TRL end, systems tended to be at the concept demonstration stage. The tests of these systems could be supervised by JVAB's less experienced personnel, since the nature of the systems being tested—and their security features—would likely change significantly prior to their release. In many respects, JVAB's value to the developers of these systems derived from the mentoring provided by the team regarding what security issues needed to be considered as the system moved forward.

High TRL systems needed to be put through a much more comprehensive battery of tests, as they were either in the field or likely to be there soon. These systems tended to be much less malleable in their design than low TRL systems. As a consequence, general suggestions from JVAB involving significant design issues were unlikely to be heeded. What was critical for these systems was a systematic investigation whose goal was to pinpoint specific security threats. This type of comprehensive evaluation needed to be supervised by JVAB's most senior personnel.

Long Term Possibilities

Many of the long term possibilities that Rohret considered for expanding the JVAB's delivery of services were dependent upon external institutional changes or technological advances. For many of these possibilities, the JVAB could not necessarily make the change through direct action. It could, however, potentially influence these external entities to encourage such changes. The immediate question therefore became how much effort should the organization expend towards driving these possibilities, given that it was already stretched with existing commitments.

Educational Reform

As Rohret had previously stated, existing cybersecurity educational programs were unsuited to JVAB's needs. Their curricula tended to be outdated and, even where current, did not focus on the delivery of skills needed for the adaptive red team player. An ideal curriculum seeking to develop red team members might contain elements such as:

1. *Developing and maintaining a persona.* Having one or more uncompromised personas would be a major asset for individuals looking to be hired on a red team. A program that helped students develop such personas and to take steps to reduce the likelihood of being outed would make them much more valuable to the workforce. Rohret recognized that this would be no easy item to incorporate into a curriculum. Aside from the obvious ethical concerns of teaching students how to misrepresent themselves online, any cookie cutter approach to teaching the material would lead to rapid recognition by the black hat community and subsequent "outing" of the students.
2. *Provide students with access to popular black hat tools.* In order to be useful immediately upon being hired, students needed to be able to detect, deflect and apply the types of malware and other tools that were in widespread use when they graduated. This objective would go hand-in-hand with the previous element, as having an uncompromised persona would greatly increase the students' ability to access such tools.
3. *Provide students with real-world cyber defense experience.* It was nearly impossible for an individual to acquire good red team skills without first having defended a system. Whether through internships or through the creation of realistic (or real) information systems likely to be subject to real world cyberattacks, the only way to acquire current defense skills was to defend against actual hackers.

Build Visibility and Demand for Specific Skills

Rohret recognized that the nature of training offered by both educational institutions and by organizations tended to be driven by the perceived need for specific skills. While the value of assessments such as those performed by JVAB was widely recognized, many organizations failed to recognize the need for such skills in-house. Rohret felt that if the value of the ART skill set was more widely visible, the workforce pipeline of individuals with those skills would grow correspondingly. There were a number of ways that such visibility might be increased:

- *Expanded involvement of a broader range of organizations in JIFX-like venues.* Much of JVAB's motivation for participating in JIFX had been to expose industry to the value of conducting systematic vulnerability assessment of technologies *before* they reached the user community. The event's after-action reports repeatedly confirmed the perception of value that resulted from JVAB's involvement. Increasing the number and scope of experiment-focused events like JIFX could serve to build demand for the ART skill set. Unfortunately, until its short term capacity

constraints were solved, it was not clear how JVAB could service such an expansion, no matter how desirable it appeared to be.

- *Building more realistic skills tests.* Despite his suspicions regarding current cybersecurity certifications, Rohret recognized that where a test existed, there would be a natural tendency to teach to it. He could envision a test—not some multiple choice instrument but rather a hands-on activity where the participant being evaluated attempted to thwart the efforts of some nefarious simulated attacker using real world tools—that would actually provide a valid assessment of the skills that JVAB was looking for. This would not only help JVAB screen applicants, it would also encourage institutions to build curricula that would prepare students for the test.

Build an autonomous system that could perform realistic attacks.

Ultimately, if JVAB's services were going to be able to expand significantly, their labor intensity would need to be reduced. Rohret could envision the development of an automated system that could search for vulnerabilities in a technology or system through attacking it in a realistic (but carefully contained!) way. The development of such a system would require considerable investment but, once constructed, could significantly leverage JVAB's existing workforce. It might also be used as a component of a tool for testing and certification, although it would necessarily require nearly continuous updating based upon real world experience if it was to remain realistic.

While Rohret felt that the JVAB team could be instrumental in developing such a system, he also recognized that such a project would require substantial resources, and would also involve a strategic redirection of the group's activities. Did it make sense to even contemplate such a large project?

The Decision

As Rohret contemplated his options, he recognized that he was also shaping the strategic direction of the firm. Effectively, he saw four distinct paths that he might follow:

1. *Focus on maintaining the status quo.* This would involve hiring as needed to replace departing staff and turning down requests to perform assessments on a regular basis, owing to capacity constraints.
2. *Organic growth.* This would mean trying to expand the JVAB team with aggressive hiring, but only doing so based upon the group's existing funding and current demands.
3. *Rapid growth.* This would involve aggressively expanding JVAB's capacity in anticipation of growing demand. This would likely require additional grants or funding and would mean that substantial effort would be required to find the right people in the highly competitive cybersecurity job market.
4. *Strategic redirection.* This would involve diverting resources towards projects that had the potential to increase the long term availability of suitable professionals, or would leverage the availability of existing professionals. The work with the education sector and/or development of automated tools would fall into this category. In addition to requiring substantial funding, this group of options would involve substantial changes in the day-to-day operations of the firm.

Rohret also wondered if he might be missing some other viable alternative. Dealing with an environment as turbulent as cyber security and cyber warfare meant that new possibilities—and threats—were continuously emerging. He knew he could not count on yesterday's solutions to today's problems.

References

- ACM Council. (1992). *ACM code of ethics and professional conduct*. ACM website. Retrieved from <https://www.acm.org/about-acm/acm-code-of-ethics-and-professional-conduct>
- JAVAB. (2012). Adaptive red team (ART) framework specifications. *Army Research Laboratory/Unique Mission Cell (ARL/UMC) Joint Vulnerability Assessment Branch (JVAB) Internal document*, October, 1-53.
- Murphy, W. F., Murphy, S. S., Buettner Jr., R. R., & Gill, T. G. (2015). Case study of a complex informing system: Joint interagency field experimentation (JIFX). *Informing Science: The International Journal of an Emerging Transdiscipline*, 18, 63-109. Retrieved from <http://www.inform.nu/Articles/Vol18/ISJv18p063-109Murphy1655.pdf>
- Suby, M., & Dickson, F. (2015). The 2015 (ISC)² global information security workforce study. *A Frost and Sullivan Whitepaper*. April. Retrieved from <http://www.boozallen.com/insights/2015/04/2015-isc2-global-information-workforce-study>

Biography



Grandon Gill is a Professor in the Information Systems and Decision Sciences department at the *University of South Florida*, where he also serves as the Academic Director of the Doctor of Business Administration program at the *Muma College of Business*. He holds a doctorate in Management Information Systems from Harvard Business School, where he also received his M.B.A. His principal research areas are the impacts of complexity on decision-making, the diffusion of academic research findings and applying the case method to STEM education. He is currently Editor-in-Chief of *Informing Science: The International Journal of an Emerging Transdiscipline* and an Editor of the *Journal of IT Education*. He is the founding editor of *Journal of IT Education: Discussion Cases*.

Exhibit 1: General Moral Imperatives from the ACM Code of Ethics and Professional Conduct

1. GENERAL MORAL IMPERATIVES.

As an ACM member I will

1.1 Contribute to society and human well-being.

This principle concerning the quality of life of all people affirms an obligation to protect fundamental human rights and to respect the diversity of all cultures. An essential aim of computing professionals is to minimize negative consequences of computing systems, including threats to health and safety. When designing or implementing systems, computing professionals must attempt to ensure that the products of their efforts will be used in socially responsible ways, will meet social needs, and will avoid harmful effects to health and welfare.

In addition to a safe social environment, human well-being includes a safe natural environment. Therefore, computing professionals who design and develop systems must be alert to, and make others aware of, any potential damage to the local or global environment.

1.2 Avoid harm to others.

"Harm" means injury or negative consequences, such as undesirable loss of information, loss of property, property damage, or unwanted environmental impacts. This principle prohibits use of computing technology in ways that result in harm to any of the following: users, the general public, employees, employers. Harmful actions include intentional destruction or modification of files and programs leading to serious loss of resources or unnecessary expenditure of human resources such as the time and effort required to purge systems of "computer viruses."

Well-intended actions, including those that accomplish assigned duties, may lead to harm unexpectedly. In such an event the responsible person or persons are obligated to undo or mitigate the negative consequences as much as possible. One way to avoid unintentional harm is to carefully consider potential impacts on all those affected by decisions made during design and implementation.

To minimize the possibility of indirectly harming others, computing professionals must minimize malfunctions by following generally accepted standards for system design and testing. Furthermore, it is often necessary to assess the social consequences of systems to project the likelihood of any serious harm to others. If system features are misrepresented to users, coworkers, or supervisors, the individual computing professional is responsible for any resulting injury.

In the work environment the computing professional has the additional obligation to report any signs of system dangers that might result in serious personal or social damage. If one's superiors do not act to curtail or mitigate such dangers, it may be necessary to "blow the whistle" to help correct the problem or reduce the risk. However, capricious or misguided reporting of violations can, itself, be harmful. Before reporting violations, all relevant aspects of the incident must be thoroughly assessed. In particular, the assessment of risk and responsibility must be credible. It is suggested that advice be sought from other computing professionals. See [principle 2.5](#) regarding thorough evaluations.

1.3 Be honest and trustworthy.

Honesty is an essential component of trust. Without trust an organization cannot function effectively. The honest computing professional will not make deliberately false or deceptive claims about a system or system design, but will instead provide full disclosure of all pertinent system limitations and problems.

A computer professional has a duty to be honest about his or her own qualifications, and about any circumstances that might lead to conflicts of interest.

Membership in volunteer organizations such as ACM may at times place individuals in situations where their statements or actions could be interpreted as carrying the "weight" of a larger group of professionals. An ACM member will exercise care to not misrepresent ACM or positions and policies of ACM or any ACM units.

1.4 Be fair and take action not to discriminate.

The values of equality, tolerance, respect for others, and the principles of equal justice govern this imperative. Discrimination on the basis of race, sex, religion, age, disability, national origin, or other such factors is an explicit violation of ACM policy and will not be tolerated.

Inequities between different groups of people may result from the use or misuse of information and technology. In a fair society, all individuals would have equal opportunity to participate in, or benefit from, the use of computer resources regardless of race, sex, religion, age, disability, national origin or other such similar factors. However, these ideals do not justify unauthorized use of computer resources nor do they provide an adequate basis for violation of any other ethical imperatives of this code.

1.5 Honor property rights including copyrights and patent.

Violation of copyrights, patents, trade secrets and the terms of license agreements is prohibited by law in most circumstances. Even when software is not so protected, such violations are contrary to professional behavior. Copies of software should be made only with proper authorization. Unauthorized duplication of materials must not be condoned.

1.6 Give proper credit for intellectual property.

Computing professionals are obligated to protect the integrity of intellectual property. Specifically, one must not take credit for other's ideas or work, even in cases where the work has not been explicitly protected by copyright, patent, etc.

1.7 Respect the privacy of others.

Computing and communication technology enables the collection and exchange of personal information on a scale unprecedented in the history of civilization. Thus there is increased potential for violating the privacy of individuals and groups. It is the responsibility of professionals to maintain the privacy and integrity of data describing individuals. This includes taking precautions to ensure the accuracy of data, as well as protecting it from unauthorized access or accidental disclosure to inappropriate individuals. Furthermore, procedures must be established to allow individuals to review their records and correct inaccuracies.

This imperative implies that only the necessary amount of personal information be collected in a system, that retention and disposal periods for that information be clearly defined and enforced, and that personal information gathered for a specific purpose not be used for other purposes without consent of the individual(s). These principles apply to electronic communications, including electronic mail, and prohibit procedures that capture or monitor electronic user data, including messages, without the permission of users or bona fide authorization related to system operation and maintenance. User data observed during the normal duties of system operation and maintenance must be treated with strictest confidentiality, except in cases where it is evidence for the violation of law, organizational regulations, or this Code. In these cases, the nature or contents of that information must be disclosed only to proper authorities.

Source: ACM Council (1992).

Exhibit 2: Adaptive Red Team Framework Methodology Defined

ARTs utilize specialized skill sets that enable them to identify, validate, demonstrate, and mitigate vulnerabilities for rapidly developed emerging, and operational technologies intended for use by the warfighter, the DOD, and other government organizations. This includes in-depth knowledge of, and experience with ancillary project requirements that may not be recognized as intricate to the technology or process being assessed. Ancillary requirements include, but are not limited to:

- Resiliency: ability to withstand or recover after an attack
- Resonance: negative effects associated with integrated components/systems
- Anticipate: using analytical data to predict, prevent and mitigate adverse effects
- Logistical requirements (shipping, notifications, orders, etc.)
- Communications prior to and during operations to include Operations Security (OPSEC) and social engineering/social network attacks
- Local (Area of Responsibility (AOR)) infrastructure
- Environmental factors
- Conflicting Concept of Operations (CONOPS) between organizations or allies

Specifically, an ART conducts research, demonstrations, and assessments to determine the processes and equipment required to successfully accomplish comprehensive vulnerability assessments on rapidly fielded and emerging technologies. Other information related capabilities (IRC) and variables outside of computer network security (CNS), radio frequency (RF) security, and electronic warfare (EW), such as cyber space operations, operations security (OPSEC), military deception (MILDEC), military information support operations (MISO), and environmental factors are included in the overall vulnerability assessment and reporting.

An ART will use adversary methods, techniques, and equipment as the baseline for researching and conducting a comprehensive vulnerability assessment. The establishment of key processes, procedures, and the assembly of a suite of assessment hardware and software, best suited to emulate the adversary and accurately identify and mitigate vulnerabilities based on the mission and the systems specific requirements, is the task of an ART.

An ART will research and conduct vulnerability assessments from the adversary's point of view; exploiting the targeted system/organization to meet documented goals. IRC vulnerabilities will be assessed in both research laboratories and simulated military/exercise environments (e.g. test ranges, military bases, and other locales that may replicate specific operational landscapes) primarily against unclassified COTS and GOTS capabilities. These assessments will be conducted in accordance with validated and anticipated adversary tactics, TTPs and/or their CONOPS. These assessments will also include new technologies and TTPs that may introduce an exploitable vulnerability into soon-to-be-fielded systems.

Source: JVAB (2012, p. 20)