Crowdsourcing Expert Performance to Improve Training at Cyber Speed

Janet Cichelli
Serco North America
Rockville, Maryland
Janet.Cichelli@serco-na.com

Maureen Baginski
Serco North America
Columbia, Maryland
Maureen.Baginski@serco-na.com

ABSTRACT

With the continued escalation in the volume, sophistication, and complexity of network attacks, the training of government professionals with the cyber skills needed to operate, maintain and defend computer networks has become increasingly urgent. In fact, maintaining a robust workforce of educated and trained "cyber-bodyguards" is a real and present national security issue. And as threats continue to change and evolve at cyber speed, training must also continue to rapidly improve and adapt to meet new and unknown threats.

The use of *crowdsourcing*—a social networking technology that leverages the "collective experience" of the group—is rapidly showing promise as a way to capture high-quality experiential knowledge and lessons learned in a more efficient way than previous methods. The human performance domain typically employs cognitive analysis and direct workplace observation to identify what experts know and do and then feed that into improvement of training methods and content. But this process can be time consuming, has logistical challenges, and provides a very limited sampling upon which to make continuous improvements in training. Applying controlled "internal crowdsourcing" techniques, expert cyber security practitioners and intelligence analysts can be immersed in a game-based environment where their performance patterns are tracked, their insights captured, and new detection patterns discerned. These findings can then be fed back—in rapid succession—to improve training scenarios and methods.

This paper presents a methodology to help meet cyberspace requirements and address the current talent gap using a rapid and agile process for continuous improvement in cyber training. The authors present a framework whereby crowdsourcing can be employed to capture human performance in a game-based network environment. The convergence of crowdsourcing and gaming provides a new and effective way for current training development models to integrate lessons learned from expert performers and improve personnel proficiency at cyber speed.

ABOUT THE AUTHORS

Ms. Janet Cichelli has over twenty-five years experience in the field of technology management and new concept development. She is a recognized solution architect and thought leader in the areas of emerging technologies, human performance consulting, and knowledge management systems. She currently serves as Chief Technologist for Serco North America, where she provides strategic consulting to Defense and Intelligence organizations in the design and implementation of technology solutions that enable a high-performing workforce.

Dr. Maureen Baginski has a distinguished back-ground with almost three decades of service in the United States Intelligence Community and currently serves as Vice President of the Intelligence business and National Security Advisor for Serco North America. Ms. Baginski previously served as the FBI's Executive Assistant Director for Intelligence where she was responsible for establishing and managing the FBI's first-ever intelligence program. She also served at the National Security Agency (NSA), where she held a variety of positions, including Signals Intelligence (SIGINT) Director, Senior Operations Officer in the National Security Operations Center; Executive Assistant to the Director of NSA/Central Security Service, Chief Officer of the Director; Assistant Deputy Director of Technology and Systems; and lead analyst for the Soviet Union.

Crowdsourcing Expert Performance to Improve Training at Cyber Speed

Janet Cichelli
Serco North America
Rockville, Maryland
Janet.Cichelli@serco-na.com

Maureen Baginski
Serco North America
Columbia, Maryland
Maureen.Baginski@serco-na.com

INTRODUCTION

Within days of taking office, the Obama Administration ordered a "clean slate" review of U.S. policies and structures for addressing the urgent challenges of cyber security. After two months of input from top cyber experts, the White House issued a comprehensive report (Cyber Policy Review, 2009) that states:

> "Cyber security risks pose some of the most serious economic and national security challenges of the 21st century… Growing arrays of state and non-state actors are compromising, stealing, changing, or destroying information and could cause critical disruptions to U.S. systems."

Some experts have warned a "cyber world war" is already underway with pervasive espionage and information extraction activity taking place across the U.S. cyber infrastructure, and reconnaissance and exploitation of critical infrastructure assets in preparation for actual attack and disabling of our critical systems. Such an attack could begin from a seemingly benign scenario:

> *The crisis began when college basketball fans downloaded a free March Madness application to their smart phones. The app hid spyware that stole passwords, intercepted e-mails and created havoc. Soon 60 million cell phones were dead. The Internet crashed, finance and commerce collapsed, and most of the nation's electric grid went dark. White House aides discussed putting the Army in American cities (Drogin, 2010).*

Scenarios such as these may sound like science fiction, but as the world becomes increasingly interconnected via the Internet, our nation's ability to safeguard critical networks, such as the electric grid, military assets, the financial sector and telecommunications, is quickly becoming a top national priority. No one knows how vulnerable we really are because the most costly attacks have not been made public. But we are probably a lot more vulnerable than we'd like to be. Attacks on government networks are also ubiquitous. According to a report by the Center for Strategic and International Studies, NASA and the departments of Defense, Homeland Security and Commerce have "all suffered major intrusions by unknown foreign entities" (Drogin 2010).

With this continued escalation in the volume, sophistication, and complexity of network attacks, the training of government professionals with the cyber skills needed to operate, maintain and defend computer networks has become increasingly urgent. Brig. Gen. John Davis, director of current operations for the Defense Department's Cyber Command, said, "There is a finite amount of human capital associated with the skill set required to do this job. There is a need for training and education for the skill set needed to do this job effectively - we need to grow the workforce" (Beasley, 2011).

In fact, maintaining a robust workforce of educated and trained "cyber-bodyguards" is considered by many to be a real and present national security issue. And as threats continue to change and evolve at cyber speed, training must also continue to rapidly improve and adapt to meet new and unknown threats.

The current demand for trained and competent cyber talent far exceeds the availability. So, to ameliorate this cyber *talent gap*, we must be able to select and develop people to higher levels of proficiency, and much more quickly. There is simply not the time to do it using a traditional, linear military training model. And cyber training must be able to continually improve to outpace an increasingly assertive and cunning adversary. The human performance domain typically employs cognitive analysis and direct workplace observation to identify what experts know and do and then feed that into improvement of training methods and content. But this process can be time consuming, has logistical challenges, and provides a very limited sampling upon which to make the continuous improvements in training. What is needed is a rapid and agile process for continuous training improvement.

Fortunately, recent developments in Web 2.0 capabilities—specifically *crowdsourcing*— are rapidly showing promise as a way to capture high-quality experiential knowledge and lessons learned that can be applied to improve training in a more efficient way.

CROWDSOURCING

Fundamentally, *crowdsourcing* is a social networking method that leverages the "collective experience" of a group. Wikipedia defines it as, "the act of outsourcing tasks, traditionally performed by an employee or contractor, to an undefined, large group of people or community (a "crowd"), through an open call" (www.wikipedia.com).

Crowdsourcing essentially brings problems or challenges directly to the masses, who serve as 21st century "free agents" to solve problems and answer compelling questions, often in innovative and out-of-the-box ways. Within the context of federal training efforts, crowdsourcing provides the ability to rapidly harness the cognitive and creative abilities of large numbers of people in a professional community to collect orders of magnitude more performance data that can then be used to identify training gaps and opportunities for improvement.

Internal Crowdsourcing

While the results from individual crowdsourcing participants can be taken and applied, the power lies in the aggregate of participants' ideas and actions. Crowdsourcing, then, by definition, is most effective with very large crowds. However, participation may need to be limited in a government cyber environment due to access restrictions to sensitive information and privileged network operations. But government organizations can limit participation and still realize many of the benefits by constraining the size and access of the community and practicing *internal* crowdsourcing.

With internal crowdsourcing, the participants can be limited to any subset of the workforce. Examples might include: personnel working within a specific organization, those in an identified job role, those identified as exceptional performers, or those with particular job experience and/or meeting training prerequisites.

Advantages

So how does internal crowdsourcing provide an advantage over the insights and experiential knowledge that an individual in an organization could provide? Simply stated, the *collective* experience, knowledge, and insight is richer than any single individual's. This is broadly acknowledged across government and is a major reason why organizations have increased their efforts to provide knowledge-sharing systems that promote teamwork and collaboration.

Effective internal crowdsourcing in the cyber environment requires a diverse group of participants, each representing different specialties and knowledge in niche areas. Cyber design, technology, and operations are far too complex for a group of experts to impart all of the knowledge and breadth of experience required to fully understand and analyze the domain and its complex systems. Complexity and diversity in the cyber domain is comparable to the medical domain, where no single practitioner can assess and treat all diseases and conditions; and therefore must rely on a large cadre of specialists and researchers. Further, the cyber domain is dynamic and rapidly changing, so that new knowledge is rarely disseminated in time for all participants to become aware of it and able to put new knowledge to use universally. Through internal crowdsourcing, both explicit and experiential knowledge can be captured efficiently from a very large and diverse population.

A recent position paper (Erickson, 2011) states three main advantages of crowdsourcing, each of which is relevant to the consideration of internal crowdsourcing for helping to provide continuous improvements in training:

1. ***Speed***: produces results more quickly by multiplying effort;
2. ***Quality***: produces higher quality results by integrating diverse input;
3. ***Legitimacy***: produces results that are more legitimate by virtue of representing a community of practice.

Speed. Some crowdsourcing systems, says Erickson, add value because they can perform a task more quickly than an individual. Classic examples he cites (2011) include von Ahn's ESP Game in which people generate textual labels for images, or Galaxy Zoo in which people classify images of galaxies as spiral or elliptical. This is usually managed by recruiting large numbers of people who perform one or two very simple tasks. Their results are then captured and integrated. The true power in this kind of crowdsourcing lies in its ability to create a situation in which many people will perform the task even though it may be trivially simple.

Quality. A second way that crowdsourcing systems add value is by producing higher quality results. Wikipedia (www.wikipedia.com) has become a classic example of using crowdsourced content authoring to create articles that are generally higher in quality than an individual might produce. Quality typically continues to increase with edits, resulting from the diversity of knowledge that its broad base of contributors bring to the writing task.

Legitimacy. The third way that internal crowdsourcing offers an advantage is that its results are based on *real performance* by *real people* in a highly *realistic environmen*t. In this way, it is strongly representative and could be deemed the best practice approach to a task or problem.

Crowdsourcing in Government

Crowdsourcing is gaining momentum within government as a means to gather public and organization ideas. The US Department of Homeland Security (DHS) has adopted the IdeaFactory crowdsourcing application to encourage brainstorming of new ideas by its employees. DHS is using the platform to encourage its employees to come forward with new ideas on how to do things. Other employees can then rate the ideas, comment on them, pick favorites, and forward them to others (Hoover, 2009).

The Open Innovation Portal, launched by the U.S. Education Department (ED), aims to use crowdsourcing to address educational challenges ranging from high school dropout rates to low reading, math, and science scores. The initiative is part of a new White House effort to encourage innovative collaboration across all industry sectors. Users also can post ideas and see, review, and rate ideas posted by others based on need, impact, evidence, innovation, and scalability (Stansbury, 2010).

But only very recently has crowdsourcing begun to emerge as a way to refine and improve training. In fact, the most tremendous promise is being shown by combining the presence and engagement of online serious games with the power of crowdsourcing methods.

CONVERGENCE OF CROWDSOURCING & GAMING

Serious games can provide for the constructive channeling of human insights and experiential knowledge, thereby providing an exceptional framework for crowdsourcing expert human performance. During game play, users will often perform tasks, solve a problem, or make a decision in a way they might not otherwise be able to articulate. But those insights—or *tacit knowledge*—can be captured and cycled back into the training improvement process for the update, validation, and enhancement of training content (e.g., scenarios) and instructional methods.

War gaming is something that military organizations do regularly in order to test their readiness and competencies in scenarios that match the real thing as much as possible. In recent years, the U.S. military has bought into gaming in a much larger way, through the use of multi-player, online *serious games*. The rationale for serious games is to attempt to capture some percent of the large amount of human brainpower expended every day in playing games and harness it toward a useful social goal or for training and readiness.

Electronic game play has become widespread for both serious and recreational use. Our society spends 3 billion hours each week playing online games (McCann, 2011). According to Luis von Ahn, a researcher at Carnegie Melon University, humans spend nine billion hours alone playing Solitaire every year (Rigby, 2009). Ahn has contrasted this huge number to the number of human-hours it took to accomplish other highly significant works in which humanity takes pride: the Empire State Building (7 million human-hours) and the Panama Canal (20 million human-hours). The results are both jaw-dropping and very intriguing related to harnessing the collective input of the crowd in serious games. If we could take the amount of voluntary human effort that was going into playing solitaire and channel it to these previously mentioned projects instead, the Empire State Building could have been completed in 6.8 hours and the Panama Canal could have been built in one day (Barquin, 2011).

In a thought-provoking TED Talk (2010), Jane McGonigal discussed the tremendous potential that could be realized if the energy and motivation found in today's gaming culture could be harnessed and directed to develop personnel proficiency in areas of national concern. She cited a Carnegie Mellon University study, which found that the average young person today in a country like the U.S. with a strong gaming culture will have spent 10,000 hours playing online games by the age of 21. This number—10,000 hours—is a very relevant number for two reasons. First, for children in the United States, 10,080 hours is the exact amount of time that each will spend in school from fifth grade to high school graduation with perfect attendance.

And the second reason it is relevant is related to Malcom Gladwell's “10,000 hour theory of success” as stated in his book Outliers (2008). This theory is based on cognitive science research that states, “If we can master 10,000 hours at effortful study, **at anything** by the age of 21, we will be virtuosos at it. In fact, we will be as good as the greatest people in the world at that particular thing.”

By combining serious games with controlled “internal crowdsourcing” techniques, cyber security practitioners can be immersed in a game-based environment where their performance patterns are tracked and new insights and detection patterns discerned. This new information is then fed back—in rapid succession—to improve training scenarios and methods. Sound far-fetched? In fact, several new training projects, combining crowdsourcing and serious games, have just begun to emerge.

The United States Navy has begun crowdsourcing tactics for fighting Somali pirates and securing the Horn of Africa, through a new video game project called MMOWGLI (Massive Multiplayer Online WarGame Leveraging the Internet). MMOWGLI was developed by the Office of Naval Research (ONR) to help solve difficult strategic problems. The MMOWGLI game, scheduled to be launched in summer 2011 focuses exclusively on combating Somalian piracy (Ungerleider, 2011).

MMOWGLI was developed by the ONR with assistance from the Institute for the Future and the Naval Postgraduate School. For the Navy, the use of massively multiplayer games to solicit military strategy is a no-brainer. Players in the game are not compensated for their time and ideas--and the military command structure has access to informed advice and strategy that they otherwise wouldn't.

In another promising effort, a crowdsourcing project called World Without Oil (WWO) is self-described as “a massively collaborative imagining” of the first 32 weeks of a global oil crisis (http://worldwithoutoil.org).

WWO is a serious game that invites people from all walks of life to contribute “collective imagination” to confront the real-world issue of declining oil resources. It is significant in its use of crowdsourcing and games as democratic, collaborative platform for exploring possible future scenarios and sparking future-changing actions.

And finally, the U.S. Army has bought into gaming in a much larger way, recently starting to incorporate crowdsourcing methods to capture soldier actions and insights. In 2010, the Army partnered with the University of Southern California Institute for Creative Technology to establish the Mobile Counter-IED Interactive Trainer, or MCIT, at Fort Bragg, N.C. to assist soldiers in recognizing and reacting to improvised explosive devices. In that time, over 15,000 soldiers nationwide have gone through the modified shipping containers featuring fictional video narratives and a computer game that are now set up at three different bases across the country (Kim, 2010).

The Mobile Counter-IED Interactive Trainer (MCIT) tries to expose soldiers to the real-life environment around making and evading IEDs. It was developed on the model of a Hollywood set and has the look and feel of the real site (Barquin, 2011). While MCIT in itself is a very effective pre-deployment training tool, the program is beginning to use crowdsourcing games to capture what returning soldiers have learned about IED detection and avoidance. These soldiers are immersed in a game where they are challenged to help predict ambush sites and IED hotspot locations. This “train first, then go down range, then give feedback” training cycle is providing valuable insights and lessons learned that the Army can use to immediately refine training for those about to deploy.

Using serious games for crowdsourcing expert performance has several benefits over other, more common, crowdsourcing methods. Chief among these is that games can offer participants a strong motivational incentive. When the game is engaging, then players will want to play, and want to continue to play, in exchange for the intrinsic enjoyment. This stands in stark contrast to other potential methods of crowsdsourcing or survey analyses, where the motivation must either be one of pure altruism or where players are compensated to participate.

APPLICABILITY TO THE CYBER INTELLIGENCE MISSION

In conjunction with the current efforts to protect and defend networks and systems in the cyber domain, there is a broader need to provide training for conducting all-source and all-method cyber *intelligence* operations (and counter-intelligence operations). If we think of the IT and cyber hardware and software infrastructure, with all of its vulnerabilities and points of entry and exploitation as the *means* to an end for our adversaries to conduct espionage and attack, we must also consider the motivations, tactics, dependencies, and desired consequences for both defending and attacking the cyber domain. The breadth of skills and tools needed to execute a comprehensive cyber strategy must go beyond the technical cyber security expertise used to manage and exploit the networks and systems.

The emerging definition of a *cyber intelligence analyst* is an almost super-human combination of cyber security expert and traditional intelligence analyst. The reality of combining human intelligence, political and cultural insight, all-source analysis, and deep cyber technical knowledge places even greater pressure on training. A game-based crowdsourcing framework can capture and manage this very dynamic set of learning requirements, and can continue to adapt as many of the rules of analysis continue to be written.

METHODOLOGY

To be operationally relevant and efficient, a crowd computing gaming platform used to solicit cyber detection and prevention strategies requires a standard framework and a process to collect, capture, analyze and use performance information. In its most basic form, the crowdsourcing framework must capture and collect desired *INPUT*, and provide aggregate and actionable *OUTPUT*.

The framework also relies on a pre-defined audience. In the cyber training context that is the focus of this paper, that audience is characterized as a large and diverse sample of accomplished performers.

The INPUT collected by the crowdsourcing framework will be in these two categories:

1) The activities and actions performed in a given situation, and
2) The crowd's opinions, decisions and insights collected on the topics and problems presented.

The OUTPUT desired is an aggregate of data/information that can be analyzed for consensus, anomalies, or performance themes.

In the Cyber Network Defense (CND) realm, players will be immersed in 3D game play where they monitor network traffic to take actions that prevent and identify sophisticated cyber threats. As the game scenario unfolds, they are required to deal with how their preventive plans failed or succeeded and must determine new strategies.

In the Cyber Operations and Intelligence (COI) realm, players will interact in a similar 3D game play, but will deal with a higher level of human, political, social, natural, and cyber activity. This is more of a cyber war-gaming exercise with critical infrastructure (food, energy, finance, water, etc.) attack events and information infrastructure espionage activities occurring in a world arena. As the game scenario unfolds players will discover, report, and respond to trends and incidents in order to maintain the integrity, trust, and availability of information and resources within the cyber domain.

Both the CND and COI realms can be merged into an all-inclusive game scenario where participants can take on any of several roles to either game at the high level exercise, or delve into the low-level cyber defense and "hacking" activities. Interaction and cooperation between these groups will be highly useful in understanding comprehensive cyber strategy in the real world.

Semi-Automated Capture of Cyber Tactics. Using a process similar to that employed in MMOWGLI (Ungerleider, 2011), a scenario will trigger a decision point: What do you do now? A text input area records the participant's brief answer. Current capabilities within the cyber intelligence domain allow for semi-automated collection of both actions and textual input from participants. The process will result in a series of datasets and hierarchical taxonomies. These datasets are built dynamically by large numbers of participants and will be used to collect general, domain-specific, and community-biased information. The resulting datasets will have a quantitative confidence assessment applied to rank the actions and feedback, as reflected by the crowd's activity.

Once player input is collected, fellow players will be able to comment or rank another player's tactic for validation. Comments can include adding additional information to expand the idea, asking a clarifying question, or challenging the original input. Then, during each round of game play, which can be executed over an extended period of time, these ideas and tactics will continue to mature and be refined.

Automated Game Play Management using Artificial Intelligence (AI). While current focus involves the semi-automated capture of performer data, research points to the future ability to include automated game play management for even greater efficiencies. Current research is being conducted around automated game play management, using AI analyzers to drive the real-time tuning and adjustment of the task, thereby modifying game parameters automatically to meet goals.

Representing Crowdsourcing Output. At the end of the planned number of rounds, the game will output the collected data and information to a repository. From that, we can display a logical treeing of the information collected. Using a hyperbolic tree view helps determine the relationship, trends, and validity of information collected, to inform decisions regarding changes in training that are needed. It is important to identify different associations and among entities, and using a hyperbolic view and a hierarchical list view allows the training team to visualize these entity relationships as a semantic tree representation of captured ideas, tactics, and information.

CONCLUSION

We now have the means to mobilize the brainpower of our nation's finest cyber defenders – using the power of serious gaming combined with simple crowdsourcing tools. Leveraging this "collective brain" of top performers in the federal cyber workforce promises to produce a higher-quality result in a more efficient way than could be produced previously.

As envisioned by the ONR, this approach of combining crowdsourcing with serious games provides for novel combinations and complex interactions of ideas that "otherwise might not emerge from more traditional wargame approaches" (Ungerleider, 2011). Serious games can incorporate crowdsourcing, to scale efforts that might otherwise be difficult for teams to do alone, or elicit information or monitor activities that people might otherwise be unable to articulate (i.e., making *tacit* knowledge *explicit*).

The bottom line is that serious games are becoming important tools for problem solving that harness crowds of users into contributing their time and knowledge in a collective way. Through a rapid and agile process, the convergence of crowdsourcing and gaming provides a new and effective way for current training development models to integrate lessons learned from expert performers and improve personnel proficiency.

ACKNOWLEDGEMENTS

The authors would like to acknowledge and thank Mr. Jim Schifalacqua, Director of the Cyber Center of Excellence at Serco North America, for his contributions and support of this paper.

REFERENCES

Barquin, Ramon. "Games and Business Intelligence." B-eye-Network blog. 3/15/11. See http://www.b-eye-network.com/view/15013.

Beasley, Meg. "DoD Wants Common Cyber Picture." Federal News Radio online. 2/24/11. See http://www.federalnewsradio.com/?nid=35&sid=2283664.

Cyber Policy Review: Assuring a Trusted and Resilient Information and Communication Infrastructure. 5/29/09. Report issued by the White House.

Drogin, Bob. "In a Doomsday Cyber Attack Scenario, Answers Are Unsettling." *Los Angeles Times* online. 2/17/10. See http://articles.latimes.com/2010/feb/17/nation/la-na-cyber-attack17-2010feb17.

Erickson, Thomas. "Some Thoughts on a Framework for Crowdsourcing: A Position Paper for the CHI 2011 Workshop on Crowdsourcing and Human Computation." CHI 2011 Workshop on Crowdsourcing and Human Computation. 5/8/11.

Gladwell, Malcolm. *Outliers: The Story of Success*. New York, NY: Little, Brown & Company. 2008.

Hoover, Nicholas J. "Homeland Security Seeks New Ideas." Information Week. 11/4/09. See http://www.informationweek.com/news/government/info-management/221600274.

Kim, Julia. "ICT at the Army Science Conference in Orlando." 12/3/10. ICT blog. See http://ict.usc.edu/blog/ict-at-the-army-science-conference-in-orlando.

McCann, Sandi. "Trends We're Watching: Social Game Marketing." The Heinrich Report blog. 6/14/11. See http://www.heinrich.com/hblog/index.php/tag/social-game-marketing.

McGonigal, Jane. "Gaming Can Make a Better World." TED Talk. Recorded March 2010. See http://www.ted.com/talks/jane_mcgonigal_gaming_can_make_a_better_world.html.

Rigby, Ben. "Information Age Volunteerism: Open Sourced! Crowdsourced!" Tech President blog. 2/6/09. See http://techpresident.com/blog-entry/information-age-volunteerism-open-sourced-crowdsourced.

Stansbury, Meris. “Feds Turn to ‘Crowdsourcing’ for Educational Innovation.” eSchool News. 5/10/10. See http://www.eschoolnews.com/2010/05/10/feds-turn-to-crowdsourcing-for-educational-innovation.

Ungerleider, Neal. “Wannabe SEALs Help U.S. Navy Hunt Pirates in Massively Multiplayer Game.” FastCompany online. 5/10/11. See http://www.fastcompany.com/1752574/the-us-navys-massively-multiplayer-pirate-hunting-game.

Wikipedia. http://www.wikipedia.com.