



Building a Bridge across the Transition Chasm

Anita D'Amico, Brianne O'Brien, and Mark Larkin | Secure Decisions, Division of Applied Visions, Inc.

An R&D team shares the techniques they used to transition government-sponsored cybersecurity research to operational environments, including ways to overcome challenges of team composition, government regulations, deployment surprises, and funding cycles.

Transitioning US government-sponsored research into operational environments is difficult and can be frustrating. Despite more than a decade of government investment in cybersecurity research and technology development, there have been relatively few successful transitions across the metaphorical chasm between government-funded research and operational use.¹ Beyond a few large players, many R&D organizations performing work for the US Department of Homeland Security (DHS), the Department of Defense (DoD), and intelligence community research agencies lack the right staff, infrastructure, tenacity, or resources to successfully conduct technology transition. A similar observation can be made about many of the research's government sponsors. Often, successful transition is due to a dedicated program manager and taking advantage of "opportunistic channels of demonstration, partnering, and occasional good fortune."²

In this article, we describe the experience behind these observations and recommend activities through which research organizations can improve the likelihood of technology transition. Our recommendations are based on lessons learned through three technology transitions, each with increasing success and each spanning many years of work with our government sponsors.

Prior work published on this topic has been largely from the perspective of either government sponsors of research or representatives of acquisition organizations (who are appropriately concerned about return on the considerable investment made in research) or from independent evaluators of prototype technologies.³

This article offers the perspective of researchers and developers who have parlayed lessons learned from prior successes and failures encountered as we've moved up the technology readiness level (TRL) ladder.⁴ Each rung of the TRL ladder reflects a step in the technology's maturation starting at TRL 1, in which basic principles have been observed and reported, and ending at TRL 9, in which a technology has been proven in real operations. Program managers and researchers assess technical progress against the TRL ladder; higher rungs represent more mature technology that is eligible for transition.

We limit this discussion to the transition of cybersecurity R&D sponsored by US government agencies. It might be tempting to view this as similar to commercial product development, but it's different in many ways. Commercial technology development is performed by R&D staff aware of and focused exclusively on the company's market, guided by professional product managers with direct access to potential users,




Figure 1. Focus and concentration of resources on the transition path, and technology readiness levels (TRLs). The figure shows the relative concentration of resources dedicated to various activities (research, requirements, development, and so forth) along the transition path. Milestone diamonds show where these three systems fit relative to TRLs.

with budgets and infrastructure for marketing, product launches, and collecting consumer feedback. By contrast, government-sponsored research is constrained by user access, limited budgets, and restrictions on information dissemination.

What Is Technology Transition?

We define technology transition as the deployment of results of cybersecurity R&D in a form and to a location in which it can be used by people or systems working in government and industry operational environments. This definition both simplifies and broadens that of James Dobbins, who writes, “Technology transition is the process by which technology deemed to be of significant use to the operational military community is transitioned from the science and technology environment to a military operational field unit for evaluation and then incorporated into an existing acquisition program or identified as the subject matter for a new acquisition program.”⁵

Other definitions of technology transition, such as transitioning a technology from one government R&D organization to another⁶ or activities such as transferring patents, publishing results, or licensing technology,⁷ are not our focus here.

As cybersecurity researchers motivated to develop

new technologies to achieve greater security, we are mindful that if the results of our work aren’t used, they’ll have little impact. Real transition occurs when a person or system uses the direct results of the research to add value to security activities. So, although we use the term “technology transition,” we don’t constrain “technology” to innovative hardware or software; it could be new, scientifically grounded methods, such as techniques that advance security skill acquisition or improve security analyst collaboration.

Problems Underlying Poor Tech Transition Rates

Before we launch into specific recommendations for what researchers can do to improve technology transition, let’s look at two problem areas that are the impetus for the recommendations that follow.

Distinct Transition Processes and Resources

Research sponsors and performers often treat transition as the end stage of a smooth continuum from research to prototype to operational system. It is not. Transition is a discontinuous process requiring different skills, funding, infrastructure, and measures of success at each stage.

Figure 1 shows three types of systems produced along the transition path: proof of concept, prototype, and production. Regardless of how you approach development—waterfall, agile, or some other way—you’ll create one or more instances of these types of system; each has different objectives and foci. A proof of concept (through TRL 3) illustrates a solution’s viability. A prototype (through TRL 6) is a requirements-driven working system subject to functional testing, demonstration, and initial stages of accreditation. A production system (through TRL 9) is a mature technology that can be operationally deployed. It must be well-tested, accreditation worthy, and accompanied by end-user services and infrastructure for training and support.

Figure 1 also shows the relative concentration of resources dedicated to various activities (research, requirements, development, and so forth) along the transition path. Milestone diamonds show where these three systems fit on the TRL ladder.

Different people, processes, and technologies are necessary in the different phases that yield these systems. Yet, we see solicitations that require research, development, and early-stage transition packaged up into one effort. For example, the DoD Small Business Innovation Research (SBIR) program (www.acq.osd.mil/osbp/sbir/solicitations/sbir20123/index.shtml) incorporates the proposer’s commercialization strategy, private investors’ funding commitments, and transition partners’ follow-on funding commitments into the evaluation of even its earliest Phase I proposals. Some SBIR Phase II topics require a demonstration of the capability “in an existing CNDSP [Computer Network Defense Service Provider] operational networking environment.”⁸ DHS requires that their “Type I—New Technologies” R&D programs include “technology demonstrations in an operational environment.”⁸

Bringing a technology into an operational environment (TRL 7) requires many of the elements of full technology transition, such as robustness, certifications, and operational champions. The implicit assumption is that the principal investigator, who has the technical expertise to achieve TRL 3, can aptly handle the engineering rigors of TRLs 4 through 9 and the relationship-building and certification hurdles required of TRLs 7 through 9. Many researchers (including us) and their sponsors have failed to recognize the importance of other skill sets in positioning a new technology for eventual deployment and, even if they recognized this, might not have had the financial resources to support the additional skill sets. The commercial security product world, by contrast, uses a product manager as the bridge between research and product launch. DHS, borrowing from the commercial model, now funds a new

category of cyber research: a “Type III—Mature Technologies” project that funds the transition process as a discrete phase of a research program, led by individuals with the requisite skills.⁸

Policy Barriers

Government acquisition, certification, and public dissemination policies form barriers to the adoption of technologies.

Discontinuities in funding within R&D programs—and between R&D and operational budgets—make building momentum toward transition difficult. Game-changing technologies by their very nature don’t have an existing base of effort and can suffer more from discontinuous funding than evolutionary research.

It can take years to obtain an Authority to Operate (ATO) to deploy any new technology on a government network. In that time, user champions and planned evaluators of the new technology might transfer, retire, or die (yes, it has taken that long). Export control paperwork spans many months and consumes research team resources outside the scope of the R&D project. International Traffic in Arms Regulations (ITAR) restrictions often placed on government-funded research make it difficult to disseminate the technology and commercialize it using the Internet.

Accreditation and certification requirements are often unclear, and even when clarified can take a year to fulfill. The National Information Assurance Partnership (NIAP) Common Criteria Evaluation and Validation Scheme (CCEVS) adds transition hurdles specific to cybersecurity by imposing specialized testing mandates that require considerable financial commitment. This results in delays in deployments, and can even derail transition activities. NIAP acknowledges that CCEVS isn’t working, in part because “evaluations are costly, not repeatable among the schemes, and take too long to meet market needs.”⁹ By the time a Common Criteria evaluation is complete, often at a cost of US\$200,000 or more, new versions of the technology have been developed, making the certified technology outdated.

These examples aren’t just speed bumps that slow down transition; they often demotivate and sap the will of the team to see the technology through its transition phase.

Recommendations: Building the Bridge to Transition

Our recommendations to other cybersecurity researchers are based on three case studies in which each system moved from TRL 1 to 9. However, the “staying power” of each transition varied from months to years.

SecureScope was our first security event visualization system, developed with SBIR funding from the US

Air Force Research Laboratory (AFRL) and DARPA.¹⁰ The path from TRL 1 to 9 spanned seven years of interrupted sponsorship, filled in by company investment. It was incorporated into an Advanced Concept Technology Demonstration and a Joint Warrior Interoperability Demonstration and licensed for use by other agencies. Its transition was cut prematurely when operational sites suddenly required the (then-new) NIAP CCEVS certification: it was pushed off the bridge and into the transition chasm by a certification requirement that didn't even exist when we started. We used

the lessons learned and technical designs from Secure-Scope to advance the transition of future systems.

VIAssist is a visual analysis platform for evaluation of network flow and security data.¹¹ VIAssist climbed the TRL ladder over six years of start-and-stop funding, during which we developed proof-of-concept, prototype, and production versions. The VIAssist foundational research, proof of concept, prototype development, and highly rated evaluation in a Coalition Warrior Interoperability Demonstration were funded through the Intelligence Advanced Research Projects Activity and AFRL. DHS Science and Technology and our own private funding supported VIAssist's maturation and customization for use in US-CERT, where it's now installed and awaiting an ATO. AFRL is expanding VIAssist for use within Air Force operational environments.

MeerCAT (Mobile Cyber Asset Tracks) is a set of integrated visualization tools to help cyber vulnerability analysts assess risks from wireless 802.11 threats and help penetration teams find exploitable access points.¹² It's our most successful transition, spanning four (intermittent) years of proof of concept and prototyping, followed by three years of licensing, operational deployment, and sustainment. Originally developed under a DARPA-funded SBIR, MeerCAT has entered the Phase III stage, with sustainment funding from the Defense Information Systems Agency and Naval Research Laboratory (NRL). It's been adopted and accredited for the DoD Information Assurance Certification and Accreditation Process (DIACAP) as part of NRL's Flying Squirrel wireless discovery suite. As part of that suite, MeerCAT has been downloaded more than 9,000 times by government users and is also available commercially through subscription, individual perpetual licensing, and enterprise licensing.

The lessons we learned from these case studies are the basis for our advice to fellow researchers. The categories of advice are ordered by relevance to the project phases,

although several apply across all phases. Our recommendations for the later phases, such as those related to testing, demonstrations, and deployment preparations, are well known in mature engineering organizations but are often dismissed by researchers trying to squeeze out as much technology as possible with their limited research dollars. Our message to fellow researchers is: if

you want to see your research used in the real world, you need to save some of your budget for those "nontechnical" tasks. Those seemingly superfluous activities are critical to technology transition.

A principal investigator who can successfully conduct research and grow a proof of concept to TRL 3 might not be equipped to lead the next phase.

People Are the Agents of Transition

The diverse capabilities of the people in your research program are as important to transition as the features and functions of the technology.

Change your research team's skill set. Change project leadership as the research effort moves through the three phases from laboratory to deployment. A principal investigator who can successfully conduct research and grow a proof of concept to TRL 3 might not be equipped to lead the next phase. Prototype development and production require engineers who are experienced in documenting and tracking requirements as well as making trade-offs among innovation, performance, and operational form-fit. Our observations in this area are further supported by the Software Engineering Institute.¹³

Add a product manager role to the research team. Model the project manager's job after commercial product managers who work with R&D staff to identify and prioritize functionality and refine features, and who have access to existing customers willing to test new technologies. This part-time team member will enhance the likelihood of transition by building knowledge of the market (government and commercial), raising awareness of competitive technologies and pricing, and forging relationships with stakeholders who will identify beta testers and potential transition targets.⁵ We first assigned a product management role to one of our staff during the MeerCAT development, and we attribute part of its transition success to her work. We've continued this practice in our subsequent research.

Create a vision document. To ensure that the research team and sponsors share a common vision, write it down. This vision document will crystalize the project's

motivation, clarify desired outcomes, and identify any major constraints or assumptions. Keep it short so that everyone on the team, including sponsors, can read it in less than 10 minutes. And get consensus.

Establish enduring relationships with stakeholders.

Regularly engage with at least two stakeholders who can serve as champions for your work from its earliest stages.⁵ It's important to note that we're referring here to individuals, not organizations: you're building personal relationships that might even span organizational boundaries. User champions will provide requirements, introduce you to other users, provide feedback on demonstrations, and serve as beta testers. Other stakeholders can explain accreditation requirements or assist with ATOs. Align yourself with at least two to mitigate the effects of losing one (for instance, to military rotation). Stay in touch even when your project is on hiatus or engaged in a long development cycle.

Manage stakeholder expectations. The words you use to describe your project's outcome affect the expectations of the research sponsors and potential test sites. If you're developing a system to assess a solution's viability and don't think it can withstand rigorous testing, refer to it as a "proof of concept." We learned this lesson when we referred to a very early version of VIAssist as a "prototype." Potential transition partners, excited by what they saw, believed that the VIAssist technology was ready for operational testing, but in fact it wasn't sufficiently mature. When it really was a fully functioning prototype, two years later, we had lost some earlier potential transition partners but gained new ones.

Expectation management goes in all directions. It's important to establish a vision and technology road map for the project early on and to share it with R&D staff, sponsors, and potential transition partners. Demonstrations of early work should be set in the context of the larger vision and road map so that the audience can see how a discrete output (software module, research finding, simulation) fits into the larger picture.

Align Requirements with Transition Phase and Value to Transition Target

Requirements—even broadly scoped ones in the earliest phases of R&D—provide goals around which to build both the technology and a shared vision of its final deployment state.

Establish distinct requirements for each phase. Each phase should have requirements and completion criteria that, if met, are likely to trigger the next phase's success. For a small proof-of-concept phase, a minimal set of requirements is needed to define the functionality

that demonstrates that the proposed approach is viable. During the prototype phase, requirements are expanded to produce a more functionally complete prototype application suitable for demonstration and including nonfunctional requirements such as performance, scalability, and stability needed for evaluation in an operationally relevant, closed environment. During production, requirements must reflect a fully deployable, stable, secure, and maintainable application, plus the infrastructure to sustain the system during operations.

Know the transition target milieu. Identify your transition targets' infrastructures and acquisition policies before choosing the operating system, software, and hardware on which you'll build. Our choice of Java for MeerCAT was driven in part by suitability of our user champions' environments. Our selection of two display surfaces in VIAssist was influenced by observations of multiscreen use in government security operations centers more than seven years ago, when multiple screens weren't ubiquitous. These lessons were built on our early SecureScope experience when we were blindsided by our transition target's need to print black and white screenshots of our 3D color displays.

Update requirements with stakeholder input. Requirements should be developed in collaboration with the stakeholders—particularly potential evaluators—during each phase. Stakeholder feedback on the proof of concept and early demonstrations, as well as the research team's growing awareness of the target milieu (operating systems, security, user expertise, culture), will significantly contribute to the next phase's requirements. Because stakeholders change, realign requirements with them at the beginning of each phase. For VIAssist, we identified requirements up front during a DoD-wide Cognitive Task Analysis preceding the initial proof of concept, and we realigned them after a series of demonstrations to potential users at the Joint Task Force for Global Network Operations. For MeerCAT, we refined requirements by frequent user evaluations during an agile development process.

Develop production technology with accreditation and compliance in mind. Assume that an accrediting body will ultimately review the new technology. If available, review the appropriate National Institute of Standards and Technology (NIST) protection profiles for security requirements. Expose the developing code base to software assurance tools to identify security risks inherent in the source code. Review the technology for compliance with US government accessibility standards, such as section 508 of the Rehabilitation Act of 1973.

Developing software in this fashion will accelerate the accreditation process.

Provide immediately recognizable value. Many research programs, by their very nature, focus on challenging, long-term needs. Try to assuage a more immediate pain point suffered by the transition target while working on those longer-term needs. When developing VIAssist, we found that network defenders wanted a streamlined method for producing watch-changeover briefings, so we added an easy-to-use report builder to accompany VIAssist's visual analytics and data querying technologies. This delivered the immediate relief of a report builder along with the more advanced visualizations that required a deeper learning curve of our users—a fair exchange, in their minds.

Never, Ever, Transition without Adequate Testing

Ever hear the expression “software testing is what users are for”? Don't believe it. Continuous testing is a key contributor to the robustness needed for real operations.

Don't skip on testing. Testing is as important as development in the final prototype and production phases. Although requirements define an application's expected functionality and behavior, testing provides a means to validate the application against these requirements. Of course, at some point, your technology might fail during an important demonstration, but comprehensive testing will minimize such occurrences and recovery time.

Test early and often. Testing isn't just an end-of-phase exercise. It should start during early prototype development and provide regular feedback on the system's health. Should the addition of new or updated functionality adversely affect some aspect of the system, this will be detected immediately instead of at the end of the phase, when the cause of the issue would be more difficult to identify and when time is invariably tight.

Build a test plan. Create a test plan and tailor it to the phase. The testing rigor and test artifacts become more challenging as one moves along the transition path. Test data, repeatable test cases, and documentation of test results that validate the system functionality are essential elements of successful transition and accreditation. Furthermore, the process of creating the test plan builds

consensus among developers, testers, and sponsors on how the target systems will be tested.

Use the test plan to collaborate with transition stakeholders. The test plan is a collaborative tool to be developed with stakeholder input and feedback. Use a test plan to clarify and consolidate expectations of stakeholders and your development staff. Involving the stakeholders at this stage of development clarifies and cements their expectations of the application.

“It's important to establish a vision and technology road map for the project early on and to share it with R&D staff, sponsors, and potential transition partners.”

Obtain test data that represents the transition environment. Obtain or create a test dataset—using either synthetic or real operational data—that will exercise and stress the system in a manner consistent with the transition targets. To test VIAssist, we created synthetic data with particular patterns that would produce predictable graphing results. This also allowed for the generation of data to exercise VIAssist's aggregation functionality and to conduct performance stress testing. To demonstrate VIAssist, we used a sanitized version of operational data we obtained from a transition partner.

Make and Sustain Memorable Interactions with Potential Transition Partners

Transition is built on a foundation of good relationships. New relationships start with a good first impression of your technology's value.

Always have something to show and tell to a potential transition partner. The idea is to have a compelling scenario-driven demonstration of your technology from which you can elicit feedback at all phases of the project. Even when only some of the research is done—or technologies developed—you can demonstrate snippets of progress against the storyboard of the project's vision and road map. Our SecureScope, VIAssist, and MeerCAT transitions were all built on relationships forged through live demonstrations of our progress provided within the context of a real-world security problem, even when the technology was in its early stages. If your project doesn't lend itself to a live demonstration, then create a prop, such as an information sheet or video that captures what your research is about and how it could provide value to the transition partner.

Make the demonstration memorable. Demonstrations typically focus either on features or on utility. Engineers and developers favor the feature-based

approach, describing one feature at a time. This is adequate for demonstrating a proof of concept but insufficient for prototype and production phases. Potential users are more interested in the utility of the overall tool rather than individual features. We've learned to craft demonstrations around real-world scenarios that are recognizable to potential users. This requires a demonstration-appropriate dataset that's sufficiently large and rich to have credibility with users. For example, to demonstrate MeerCAT, we collected wireless data by "wardriving" around a local shopping mall.

This illustrated the technology's value against a background of commercial enterprises handling credit cards and medical facilities handling private patient data.

Don't lose a transition opportunity because you've bored the demonstration's audience into an altered state of consciousness. Ensure that the presenter is trained to speak in public and field questions. Even a demonstration that takes a deep dive into the technology can be delivered with clarity and enthusiasm.

Go where the users are. Put your technology where your users can encounter it. Provide information sessions and demonstrations at conferences where users gather. If possible, propose the technology for inclusion in a military or homeland security exercise. While there, solicit feedback and additional requirements from subject matter experts and identify beta testers and early adoption sites. Provide printed and electronic materials that describe the technology features and benefits, the needs the technology meets, and the operational users who are most apt to benefit from the technology.

Remove barriers to global dissemination of information about the technology. The Department of State ITAR and the Department of Commerce's Export Administration Regulations laws prohibit unlicensed export of much information related to military and commercial technologies. Assess the need for a license, which is dependent on the technology's characteristics, destination, end user, and end use.¹⁴ Determine if the technology falls under export control and address it approximately one year before planned release of technical information outside the government.

Get testimonials from beta testers. Happy beta testers, particularly in government, might be reluctant to provide testimonials to the success of a prototype because of approval hurdles. Apply for that approval as early as

possible; you may get lucky. When possible, seek commercial beta testers who your transition targets recognize as knowledgeable in the domain. Industry testers with service offerings related to your technology are more apt to publicly share their views if their assessment of leading-edge technology reflects their interest in innovation and links back to their own public persona.

'Drive-by fielding' occurs when the research team drops the technology at the operational site and declares victory.

Name your program based on your transition target. Give stakeholders a way of remembering your research and associating it with what it does and how

it meets their needs. Naming

can help serve that purpose. We've failed miserably and succeeded wildly in this area. We failed by naming our cyber mission assurance technology "Camus": it was an acronym for mapping cyberassets to missions and users, and it invoked Albert Camus. Wrong! Few of our stakeholders "got it," and our self-indulgence did nothing to advance its transition. Among our successes, we renamed our wireless security research "MeerCAT" for its association with the emerging Flying Squirrel wireless security suite, and for its suggestion of "whacking" unauthorized access points as they "pop up." Everybody gets it and remembers it.

Be Prepared for Deployment and Unforeseen Obstacles

Murphy's law says that "Anything that can go wrong, will." We have all experienced the seemingly solid demo or deployment that fails at the worst possible moment. There are things you can do to prepare for and mitigate the effects of those unforeseen circumstances.

Have a deployment plan. Smooth installation requires deployment planning by both developers and the deployment site. Create a deployment plan and share it with the site well in advance of actual deployment. A typical deployment plan includes specifications for target systems (OS, platform, CPU, RAM, graphics, Security Technical Implementation Guide [STIG] configurations), external interfaces and data sources, application installer, installation instructions, and level of access needed to install an application or system. Be sure to include logistical support information, such as requirements for getting into the building; even though your technology is unclassified, you might need security clearances or escorts for uncleared personnel at some government installations.

The plan should also stipulate access requirements for people with administrative rights to the platform on

which your application will be installed as well as people with administrative rights to the external resource to which you'll be connecting (such as a database). More details of our experience with this while deploying VIAssist to US-CERT can be found elsewhere.¹⁵

Budget more time than you think it will take. It has always taken us more time than we—or our transition partners—initially estimated for installation and deployment. While the concept of installing an application at an operational site seems straightforward, many things can pop up and make this task more difficult than expected, such as improperly configured hardware, inappropriate credentials, and lack of access to site administrators.

Establish a risk management plan. A risk management plan quantifies the criticality of each risk in terms of time, budget, and staffing and offers corrective or alternative actions. Establish a risk management plan with your transition partner to address factors that could impact the deployment schedule. Examples of risks are delays in hardware and software delivery, personnel turnover, and delays in receiving security clearances.

Keep your eye on that ATO. Know your transition partner's requirements for attaining an ATO that will permit the installation of the technology on a selected government network. Going through the ATO process takes substantial resources and calendar time, so start it during the prototype phase. Because ATOs expire, remember to recertify.

ATOs are granted by Designated Approving Authorities whose requirements vary for certifying and accrediting application security. Government organizations are mandated, by Federal Information Security Management Act (FISMA) regulations, to develop and implement programs that provide security for their systems and information. Each has some latitude to direct how FISMA compliance will be achieved: compliance with certain specifications might be required, such as the NIST Risk Management Framework, DoD DIACAP, or other organization-specific requirements.

Set Up an Infrastructure for Deployment Support and Sustainment

"Drive-by fielding" occurs when the research team drops the technology at the operational site and declares victory. Although this might allow you to check the deployment box on your funding agreement, it doesn't fulfill the spirit of technology transition and might leave a bad last impression of an otherwise good program. To avoid this, give your transition sites the following types of support.

Develop user manuals. Every evaluator and operational

end user should be supplied with a user manual. Polished user manuals send a message to the evaluator that this project has progressed well past the research stages and that you've given consideration to operational needs. We have found that two levels of information are needed: a quick-start guide and detailed documentation for advanced users.

Provide end-user training. Training is needed any time the application is exposed to new users who might evaluate it, not just at final deployment. Users must be trained on how to use the technology before any evaluation. Lack of training can result in a poor evaluation of the technology and can become an impediment to the project's transition to an operational environment.

Develop training manuals and briefings as soon as the first prototype is ready. This also prepares your staff for how to present the application to new users. VIAssist training materials went through several iterations and were augmented with hands-on exercises that engaged users to learn both the fundamentals and advanced use of the system.

Set up a tech support infrastructure. Establish a technical support infrastructure to include a website with FAQs, knowledge base and documentation, and phone and email support. Hot fixes and patches should be downloadable from the Web or FTP sites. For MeerCAT, we structured different levels of support. For commercial users, we provide everything from basic help desk support to in-depth technical assistance. For the thousands of DoD Flying Squirrel users who use MeerCAT-FS, we provide behind-the-scenes technical assistance to the Flying Squirrel help desk as well as training courses to help them field user questions.

Incorporate deployment feedback into technology refinements. Set up a mechanism for continuously engaging end users for feedback—in the form of an online trouble ticket system, user forums, blogs, and webinars—with a goal of identifying new features and future developmental initiatives. Establish a plan for releasing new versions of the technology that incorporate enhancements and features that reflect the deployment site's operational needs. This will help earn a place in the transition partner's sustainment budget.

Fill the Funding Gaps

The "valley of death" aptly describes the funding gap between a TRL 6 technology demonstration and the higher rungs of the TRL ladder that lead to operational deployment. To build a bridge across the valley requires belief in the technology's value proposition and a willingness to financially support that belief.

Be prepared to fill in the funding gaps. We haven't encountered a research program that was sufficiently funded to perform research, develop technology, and transition. Yet keeping the technical expertise of the research team focused on the effort and keeping the results in front of potential transition targets are important for technology transition. You must be prepared to self-fund or seek private investment to fill in the financial gaps between research grants and fiscal years.

Keep the contract open for additional investments. With an open contract vehicle, technology champions from other agencies can transfer funds into the research grant or contract. A small amount can help fill the funding gaps between phases. Work with your program manager and contracting officer to extend the contract in anticipation of plus-ups. If you're working under an SBIR, outside investments might be eligible for matching funds, doubling the impact of an interagency funds transfer or a private investor.

"Build it and they will come" doesn't apply to cybersecurity R&D. You need the right people, the personal characteristics, tolerance for ambiguity, adequate funding, and time to succeed in technology transition.

People—your research team and the government champions—are as important to successful transition as the technology itself. Your R&D team needs researchers, engineers, testers, a usability expert (for human-mediated systems), a product manager, a legal consultant, and a project manager to deliver usable technology and get it past the ATO, certification, and export control hurdles. You should find and form relationships with anyone in the transition decision space—for instance, sponsors, potential users, evaluators, certifiers, accreditors, network administrators, and contracting officers.

Your team must possess will, tenacity, patience, and interpersonal skills to cross the transition chasm. Researchers and developers must overcome their inclination to stay in the office, engaged in intellectual pursuit, and must venture out to meet with stakeholders. Emails and phone calls have to be sent and re-sent, often over months, to get answers to simple questions about certification requirements, ATO status, when evaluators are available for training, and so on.

There's nothing clear or straightforward about the transition path, no checklist that ensures success if followed. The decision makers and policies change along the way. The R&D team has to be able to tolerate ambiguity and engage in activities with uncertain procedures and outcomes. The way forward is shown

by a series of government stakeholders (and their contract support staff) who take the time to make a call or clarify a requirement.

Continued funding is critical, but it doesn't have to come from a single source. A transition path can be traversed through a patchwork of funding sources.

Technology transition isn't fun, but it is rewarding. The reward is seeing your new technology in use, making an impact on our cybersecurity posture. ■

References

1. W.D. Maughan, "Crossing the 'Valley of Death': Transitioning Research into Commercial Products—A Personal Perspective," *Proc. IEEE Symp. Security and Privacy*, IEEE CS, 2010, pp. 21–26.
2. "A Roadmap for Cybersecurity Research," Dept. Homeland Security, Nov. 2009, page B2; www.cyber.st.dhs.gov/docs/DHS-Cybersecurity-Roadmap.pdf.
3. S.C. Paladino and J.E. Fingerman, "Cybersecurity Technology Transition: A Practical Approach," *Conf. Homeland Security—Cybersecurity Applications & Technology (CATCH 09)*, IEEE, 2009, pp. 325–330.
4. *Department of Homeland Security Science and Technology Readiness Level Calculator* (version 1.1), Final Report and User's Manual, Homeland Security Inst., 30 Sept. 2009; www.homelandsecurity.org/docs/reports/DHS_ST_RL_Calculator_report20091020.pdf.
5. J. Dobbins, "Planning for Technology Transition," *Defense AT&L*, Mar.–Apr. 2004, pp. 14–17.
6. "Manager's Guide to Technology Transition in an Evolutionary Acquisition Environment Version 2.0," Dept. Defense, Defense Acquisition Univ., June 2005.
7. A. Gillespie et al., "Key Findings and Recommendations for Technology Transfer at the ITS JPO," publication no. FHWA-JPO-11-085, Nat'l Technical Information Service, 21 Mar. 2011; http://ntl.bts.gov/lib/42000/42100/42107/FHWA-JPO-11-085_Key_Findings_Recommendations_for_Tech_Transfer_at_ITS_JPO_PDF_508.pdf.
8. 2011 DHS S&T Cyber Security Research and Development Broad Agency Announcement 11-02, amendment 00014, Dept. Homeland Security, 30 June 2011; www.fbo.gov/index?s=opportunity&mode=form&id=40161dd972cd60642ecaaa955e247067&tab=core&_cview=1.
9. "Reforming the Use of Common Criteria," National Information Assurance Partnership Common Criteria Evaluation Validation Scheme, 28 July 2012; www.niap-ccevs.org/evolution/announcements/NIA-CCEVS_Brochure_Pamphlet.pdf.
10. A. D'Amico and M. Larkin, "Methods of Visualizing Temporal Patterns in and Mission Impact of Computer Security Breaches," *Proc. DARPA Information Survivability Conference and Exposition (DISCEX 01)*, IEEE CS, 2001, pp. 343–351.
11. A. D'Amico et al., "Visual Discovery in Computer Net-

work Defense," *IEEE Computer Graphics and Applications*, vol. 27, no. 5, 2007, pp. 20–27.

12. K. Prole et al., "Wireless Cyber Assets Discovery Visualization," *Proc. 5th Int'l Workshop Visualization for Computer Security (VizSec 08)*, Springer, 2008, pp. 135–143.
13. E. Forrester, "A Lifecycle Approach to Technology Transition," Software Engineering Institute, 1 Sept. 2003; www.sei.cmu.edu/library/abstracts/news-at-sei/feature43q03.cfm.
14. "Introduction to Commerce Department Export Controls," US Dept. Commerce Bureau of Industry and Security, 5 May 2003; www.bis.doc.gov/licensing/exportingbasics.htm.
15. B. O'Brien, A. D'Amico, and M. Larkin, "Technology Transition of Network Defense Visual Analytics: Lessons Learned from Case Studies," *IEEE Int'l Conf. Technologies for Homeland Security (HST 01)*, 2011, pp. 481–486.

Anita D'Amico is director of Secure Decisions, Division of Applied Visions, Inc. Her interests include information visualization, cognitive analysis, cyber mission assurance, and technology transition to the operational environment. She holds a patent (no. 6,906,709) for techniques in network security event visualization. She received a PhD in psychology from

Adelphi University. D'Amico is a member of IEEE, the American Psychological Association, the Human Factors and Ergonomics Society, the Armed Forces Communications Electronics Association, and the Information Systems Security Association. Contact her at anita.damico@securedecisions.com.

Brianne O'Brien is a project manager in Secure Decisions, Division of Applied Visions, Inc. Her interests include software development, system testing and evaluation, accreditation and certification, cyber range development, visualization, project management, and training. O'Brien received an MBA from Adelphi University. Contact her at brianne.obrien@securedecisions.com.

Mark Larkin was a project engineer in Secure Decisions, Division of Applied Visions, Inc., and is currently with North Atlantic Industries. His interests include software development, software testing and evaluation, embedded systems, and visualization methods. He holds a patent (no. 6,906,709) for techniques in network security event visualization. Larkin received an MS in computer science from the New York Institute of Technology. Contact him at mlarkin@naii.com.

ICSE 2013

International Conference on Software Engineering

Jointly sponsored by IEEE Computer Society and ACM SIGSOFT

18 – 26 May 2013 • San Francisco, California

Early Registration until 14 April 2013

Don't miss this really special ICSE!

Great keynote speakers

Cutting-edge technical research track

Thirty-something workshops on a broad variety of topics

14 tutorials

Software Engineering in Practice track

And much, much more!



<http://2013.icse-conferences.org>