UNDERSTANDING MISSION-DRIVEN RESILIENCY WORKSHOP

Monday, March 18, 2019

MIT Beaver Works, Cambridge, MA

Authors/Editors:
Jeremy Mineweaser
Rob Lychev
Orton Huang
Martine Kalke
Reed Porada

LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

 This material is based upon work supported by the Department of the Air Force under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Department of the Air Force.

About the Workshop

MIT Lincoln Laboratory hosted an invitation-only, one-day interdisciplinary workshop entitled "Understanding Mission-Driven Resiliency" on behalf of the US Air Force, on March 18, 2019 at MIT Lincoln Laboratory Beaver Works in Cambridge, MA. Participants began to bridge the gap between government and industry to improve the resiliency of government systems to cyber attacks. The workshop focused on understanding and defining resiliency from different perspectives and included five panels devoted to discussing how different industries view and manage resiliency within their organizations, the sources of resiliency within organizations and software-intensive systems, measuring resiliency, and building resiliency within an organization or technology stack.

Background

Adversaries are investing in cyber operations capabilities that pose grave danger to the systems we rely on. At the same time, our software and computing systems are increasing in complexity and becoming less predictable. Each new device has more transistors, and more devices are in use, with ever-greater connectivity and additional levels of abstraction. Few users know exactly how their devices work; each developer understands only a fraction of the whole, yet their choices have impacts far beyond their understanding. The situation is further exacerbated by the intrinsically inseparable social and technical aspects of the problem; cyber adversaries leverage both technical and social vulnerabilities to penetrate and compromise commercial and non-profit critical infrastructure. A commercial cybersecurity market has emerged to exploit this situation, offering temporary assistance (for a fee) without delivering long-term success or permanently deterring attackers. Meanwhile, academic research continues according to its own agenda. These activities, however, are not adequately improving our understanding of system resiliency or producing secure systems up front. Many of the current approaches to managing risk from current and future cyber threats are too complicated and/or too abstract to be effective in practice. Without a shared, nuanced understanding of dependencies, mission owners cannot effectively allocate limited resources to improve resiliency. This understanding can be surprisingly difficult to achieve for large, highly distributed systems in which most stakeholders remain unaware of each other.

Workshop Schedule and Participation

Five moderated panel discussions explored an inter-related set of themes and questions:

- **Dynamics of Resiliency for the Federal Government.** Software-intensive systems within the government frequently exhibit fragility, even in benign conditions, due in part to shortfalls early in the system lifecycle. In the absence of well-crafted requirements, a wide gap may emerge between stakeholders' early expectations and their lived experience with fielded systems. This panel focused on understanding what resiliency means for the government and the challenges that the government faces when trying to achieve mission resiliency across the spectrum of conflict. It has proven challenging to acquire resilient systems, and the government often finds itself using acquired tools and services to improve resiliency of legacy systems. What are the challenges and possible paths forward given that system capabilities and complexities are ever increasing, dependencies on industry and external entities are often poorly understood, and the government levers to improve the situation appear limited?
- **Dynamics of Resiliency within Various Industries.** Companies in software-intensive industries have also been tackling resiliency challenges, but government and industry perform different missions under different constraints. What does resiliency mean for various industries at the technical and organizational levels? How do organizations identify and mitigate risks? What are the relevant tradeoffs, and are there risks too large to mitigate? How do enablers of resiliency

align across people, process, and technology? This panel was dedicated to discussing what it means for commercial entities to achieve resiliency for their missions at the technical and organizational levels, and how their methods could be leveraged by the government.

- **Sources of Resiliency.** When software-intensive systems fail to meet expectations, the blame is often placed on faulty technology, but social and technical aspects of the problem are intrinsically inseparable: the operators and maintainers of technologies are at least as important as their designers and builders. In addition, many government organizations now find themselves reliant on low-cost, feature-rich systems and services, provided by industry, that introduce additional dependencies and dynamics that are not well understood. To what extent should system owners/operators allocate resources to measuring and understanding their adversaries' capabilities and intentions? This panel was dedicated to exploring the different sources of resiliency internal and external to an organization while relying on a multitude of resiliency enablers across people, process, and technology.
- **Measuring and Characterizing Resiliency.** What are the best ways to measure and characterize resiliency? Gaining adequate insight into a system's resiliency requires conscious, sustained effort; complying with general policies or implementing standard security controls does not ensure success. Neither does a "measure everything" approach lead to positive results – instead, operators are likely to be overwhelmed during a crisis, unable to navigate effectively. Which aspects of resiliency are measurable, which are not, and why? What are the fundamental limitations? How does one aggregate qualitative and quantitative characterizations together? This panel was dedicated to exploring possible ways and limitations to measuring and characterizing resiliency quantitatively and qualitatively.
- **Building Resiliency.** This panel focused on how we can better equip our organizations to conceive, build, operate, and maintain more resilient systems. What considerations, constructs, and processes can be used to improve how we design, build, and operate our systems? How can we educate and incentivize builders, maintainers, and operators at the technological and organizational levels in order to increase performance while operating in a contested environment? How does an organization create capacity?

More than 60 people attended the workshop, including participants from government, industry, and academia. In addition to the panel discussions, the workshop offered opportunities for professional networking and small-group discussions on topics of shared interest. This summary of the workshop is being distributed to workshop participants in June 2019. The workshop raised many questions to inspire and guide future work. In a few cases, participants reached rough consensus on specific principles and practices. These are highlighted in the findings below, along with open questions to be revisited in future workshops.

Summary of Workshop Findings

The Appeal and Perils of Complex Systems

When technological advances enable new functionality, the resulting systems enjoy widespread appeal, despite their increasing complexity. Over time, a new baseline expectation is established. Our ability to create these systems now exceeds our capacity to understand them and predict their behavior and performance. In a contested environment, this presents adversaries with many opportunities for

exploitation. Defenders commonly take steps to increase robustness[1], but this addresses only the attacks and failure modes they can anticipate and model accurately. Such reductionist approaches are unlikely to succeed in the long run, because our prevailing practices for the engineering of large-scale, software-intensive, socio-technical systems cannot account for all risks to mission operations in an uncertain environment. Designers should look beyond internal technological choices and explore opportunities for operational and policy innovations, especially those that reach beyond their own organizational boundaries. Preference should be given to innovations that can be rapidly scaled up to all stakeholders without incurring runaway costs. Rather than waiting for a single, technically 'perfect' solution to emerge, decisionmakers should pursue multiple, competing approaches and set time-targeted thresholds for achieving incremental improvement.

Open Questions: How can we improve design-time prediction of the relationships between functionality, complexity, and resiliency? How can we "change the shape of the curves" that characterize these relationships?

Resiliency as a First-Class System Property

In an uncertain operational environment, long-term commitment to a static engineering design may prove difficult to defend. Such risks should be hedged. This means valuing security and resiliency as first-class properties while accepting that inconvenience and inefficiency are firmly associated with operational excellence in a contested environment. Missions and systems should be structured to accommodate changes in both technology and behavior. Modular design is not a panacea, as it depends very much on choosing the right specifications for each component. Instead, organizations should examine the overall economic landscape, beyond individual instances of technology use. Consider organizational structures and operational designs that lower the cost of control. Explore ways to deepen understanding and reduce hidden complexity. When a particular problem seems insoluble, reframe it more broadly[2,3]. Taken together, these practices may lead to more dynamic system designs that retain greater capability when attacked and enable operators to adapt quickly and effectively when new threats emerge.

Open Questions: How can we encourage technologists to broaden their horizons and seek improvements across the entirety of the socio-technical economic systems in which they are embedded? Which aspects

---

[1] A robust system is one that achieves its performance objectives while subject to specific disturbances, within known limits, based on a well-defined model of anticipated operating conditions.

[2] General Dwight Eisenhower: "Whenever I run into a problem I can't solve, I always make it bigger. I can never solve it by trying to make it smaller, but if I make it big enough, I can begin to see the outlines of a solution."

[3] Steve Landsburg, reflecting on the work of Alexander Grothendieck: "Imagine a clockmaker, who somehow has been oblivious all his life to many of the simple rules of physics. One day he accidentally drops a clock, which, to his surprise, falls to the ground. Curious, he tries it again—this time on purpose. He drops another clock. It falls to the ground. And another. What is it about clocks, he wonders, that makes them fall to the ground? He had thought he'd understood quite a bit about the workings of clocks, but apparently, he doesn't understand them quite as well as he thought he did, because he's quite unable to explain this whole falling thing. So, he plunges himself into a deeper study of the minutiae of gears, springs and winding mechanisms, looking for the key feature that causes clocks to fall. It should go without saying that our clockmaker is on the wrong track. A better strategy, for this problem anyway, would be to forget all about the inner workings of clocks and ask "What else falls when you drop it?" A little observation will then reveal that the answer is "pretty much everything", or better yet "everything that's heavier than air." Armed with this knowledge, our clockmaker is poised to discover something about the laws of gravity."

of these systems will be targeted by malicious actors, and how do their choices affect the course of system evolution?

Resiliency Is People

Resiliency has two parts: a system's reliability (or risk) profile corresponds to a characterization of its behaviors under well-defined conditions; a system's adaptability is a characterization of its ability to continuously learn from observations and improve its reliability (or risk) profile. Resiliency can only be achieved when stakeholders work together with common purpose. No isolated defender can succeed indefinitely against a determined and skilled attacker; operators must build mutual trust and be ready to assist others. Attackers will exploit technical flaws as well as organizational dysfunction, including gaps in institutional knowledge. Technical "fixes" are rarely effective at resolving dysfunction within teams. Instead, problematic behaviors must be confronted directly, on their own terms. In some cases, reorganizing how work gets done can dramatically improve outcomes, both in daily practice and when operators are responding to an attack. These kinds of improvements can be more cost effective than making major technical changes. When stakeholders for all aspects of the mission dedicate themselves to the daily practice of resiliency, everyone will be more successful at completing the mission in a contested environment. Advocates for resiliency should find ways to work with, not against, the developers who create new functionality in pursuit of mission objectives. Instead of burdening operators with additional 'security' responsibilities, empower them with the authority to change procedures when disruptions occur. Responsibility for defining (and modifying) procedures should reside with those who are accountable for mission outcomes.

Individuals may not fully realize how their fear of failure can undermine system resiliency by inhibiting opportunities for organizational learning and improvement. Too often, developers and operators face uncertain threats without the benefit of confidence-building experience performing their mission in a contested environment. Training and mission rehearsals focus on core competencies in benign cyber environments. Untested incident response procedures may become stale while adversaries develop new attacks in secret. An alternative approach would embrace and promote antifragility: system stakeholders would encounter stressing conditions in a controlled manner, then devise new procedures (or technologies) to maximize their collective capacity during actual crises. Over time, they could expand their performance limits, reduce the occurrence of operational surprises, and become more adept at creatively responding to novel situations. Adopting this approach requires collective action by stakeholders across the entire system development life cycle.

Open Question: How can we change the culture so that stakeholders appreciate the benefits of antifragility?

Real Options Enable Maneuver Amid Uncertainty

Development and operations should be organized for explicit, deliberate choices. Align each choice with the associated benefits and losses. Every stakeholder should take responsibility for their choices and understand what each asset brings to the mission. Through deliberate practice[4], much unnecessary complexity can be avoided, particularly during early phases of the lifecycle. However, the future operational environment remains uncertain, and some choices will look poor in retrospect. To ensure successful mission outcomes, organize for real competition, making multiple bets across a truly diverse array of system designs. Establish multiple, independent ways to accomplish the mission. Pursue

[4] https://jamesclear.com/deliberate-practice-theory

organizational and process innovations as eagerly as new technologies are examined. Ensure that mission plans include enough "slack" time for stakeholders to adapt to adversary actions. Always gather feedback on mission performance and convey that information to those whose decisions influenced the outcome. Cultivate continuous learning about mission-system dependencies.

Open Question: How can we promote long-term thinking in organizations where careers are built on short-term successes and decision-makers are rarely in place long enough to see the full consequences of their choices?

Risks Are Shared

All stakeholders for a given mission must work together to create strategic coherence. A simple strategy is best. Seek broad consensus on desired outcomes, then enable stakeholders to pursue competing approaches to achieving those outcomes. Once equipped with a shared narrative about resiliency, each decision-maker will be sensitive to the potential downside risks of their local adaptations. Together, they can examine how risk accumulates from individual choices to pose major challenges for all of society. However, the clear benefits of coherence must not lead to undue demands for "unity" that result in the suppression of critical or dissenting perspectives. Instead, stakeholders should adopt an explicitly experimental approach, creating real options so that missions can succeed even when major planning assumptions turn out to be invalid.

Open Questions: How can we improve collaboration on risk assessment to better identify systemic issues and take corrective steps as a society? Where should we challenge adversaries' domination of international processes for standardization of key technologies?

Experimentation Enables Collaborative Discovery

Operational plans should acknowledge the likelihood of adversarial action against software-intensive systems. While inputs may be more easily measured, operational outcomes should be the focus of continuous improvement efforts. Those efforts should be organized around incident response plans that are created, maintained, and regularly exercised by a joint collaborative team of mission stakeholders. A small number of plans should address the most likely and most dangerous contingencies. Manage residual uncertainty through a campaign of experimentation in which each plan enables whole-system discovery (by trial and error) to identify challenging scenarios and select high-leverage opportunities for improving performance. Notable measures include time to recover and time to change system configuration.

Open Question: How can we motivate stakeholders to exercise their plans regularly?

Plan for Action, Community Discussion, and Future Workshops

The workshop sponsor will apply the insights summarized in these proceedings to its ongoing program activities, specifically including its resiliency-focused projects with MIT Lincoln Laboratory.

Many participants expressed a strong desire to continue the workshop series, with subsequent events aimed at investigating the open questions and delivering products for use by the community. If you would like to participate in a future workshop, please contact the organizers to discuss your proposal. If

you know others who could contribute to the success of future workshops, please share this report with them and encourage them to contact the organizers.

About MIT Lincoln Laboratory

MIT Lincoln Laboratory is a United States Department of Defense research and development center working on problems pertinent to national security on behalf of the military services, the Office of the Secretary of Defense, and other government agencies. The Secure Resilient Systems and Technology Group focuses on the development and prototyping of new technologies and capabilities for ensuring the security and resiliency of next-generation mission-critical systems from drones and satellites, to handheld devices and miniature sensors, to high-performance secure cloud computing, to many others. Program activities consist of computer scientists, software, hardware, and electrical engineers, cryptographers, system analysts, and security architects working collaboratively on fundamental investigations, through simulation and analysis as well as design and field-testing of prototype systems.