


Federal Register Notice 88 FR 7999, [Federal Register: Request for Information on the 2023 Federal Cybersecurity Research and Development Strategic Plan](#), March 3rd, 2023

Request for Information on the 2023 Federal Cybersecurity Research and Development Strategic Plan

MITRE

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views or opinions of the U.S. Government. We bear no responsibility for the accuracy, legality, or content of the responses and external links included in this document.



MITRE's Response to the OSTP RFI Supporting the 2023 Federal Cybersecurity R&D Strategic Plan

March 3, 2023

For additional information about this response, please contact:

Duane Blackburn
Center for Data-Driven Policy
The MITRE Corporation
7596 Colshire Drive
McLean, VA 22102-7539

policy@mitre.org



About MITRE

MITRE is a not-for-profit company that works in the public interest to tackle difficult problems that challenge the safety, stability, security, and well-being of our nation. We operate multiple federally funded research and development centers (FFRDCs); participate in and lead public-private partnerships across national security and civilian agency missions; and maintain an independent technology research program in areas such as artificial intelligence, intuitive data science, quantum information science, health informatics, policy and economic expertise, trustworthy autonomy, cyber threat sharing, and cyber resiliency. MITRE's 10,000-plus employees work in the public interest to solve problems for a safer world, with scientific integrity being fundamental to our existence. We are prohibited from lobbying, do not develop or sell products, have no owners or shareholders, and do not compete with industry. Our multidisciplinary teams (including engineers, scientists, data analysts, organizational change specialists, policy professionals, and more) are thus free to dig into problems from all angles, with no political or commercial pressures to influence our decision making, technical findings, or policy recommendations.

MITRE has been at the forefront of cyber defense since the very beginning. MITRE draws from a wealth of deep technical expertise to create innovative solutions that address the ever-evolving challenges in cybersecurity.¹ We advocate a multi-faceted, interactive approach—which, in turn, broadens our impact through the power of collaboration. We know that working in partnership is crucial to national security, critical infrastructure, economic stability, and personal privacy. We serve as a trusted adviser across government (including both the national security community and federal agencies that serve citizens) and with other partners.

As part of our cybersecurity research in the public interest, MITRE has a 50-plus-year history of developing standards and tools used by the broad cybersecurity community. With frameworks like ATT&CK[®],² Engage[™],³ D3FEND[™],⁴ and CALDERA[™],⁵ as well as a host of other cybersecurity tools, MITRE arms the worldwide community of cyber defenders. We give them vital information to thwart network intruders, build resiliency against future attacks, and develop assurance to overcome possible vulnerabilities.

Since 2014, MITRE has operated the country's only FFRDC dedicated to cybersecurity and the advancement of secure technologies, which supports the National Institute of Standards and Technology's (NIST's) National Cybersecurity Center of Excellence—a state-of-the-art collaborative hub where government, industry, and academia are building practical solutions to meet the real challenges businesses face each day. The National Cybersecurity FFRDC (NCF) is focused on the serious and growing risk cyber attacks pose to economic prosperity, public safety, and national security. The NCF brings multidisciplinary teams of cybersecurity architects,

¹ Cybersecurity. 2023. MITRE, <https://www.mitre.org/focus-areas/cybersecurity>. Last accessed March 1, 2023.

² ATT&CK[®]. 2022. MITRE, <https://attack.mitre.org/>. Last accessed March 1, 2023.

³ Engage for Defenders. 2022. MITRE, <https://engage.mitre.org/defenders/>. Last accessed March 1, 2023.

⁴ D3FEND[™]. 2022. MITRE, <https://d3fend.mitre.org/>. Last accessed March 1, 2023.

⁵ CALDERA[™]-A Scalable, Automated Adversary Emulation Platform. 2021. MITRE, <https://caldera.mitre.org/>. Last accessed March 1, 2023.

engineers, social scientists, and communications professionals together with NIST to design and build usable, real-world solutions.

A more detailed discussion of MITRE's cybersecurity history and impact is provided in Appendix A. Our recommendations in this response are based on insights gained through these extensive activities.

Introduction and Overarching Recommendations

The 2019 Federal Cybersecurity Research and Development Strategic Plan remains largely accurate and viable today despite technological advancements; COVID-19-driven changes in work, education, and communication paradigms; and new global conflicts and international competition that have taken place over the ensuing four years. This is a testament that the challenges in the prior Strategic Plan were both properly identified and difficult to overcome. That said, an updated 2023 Strategic Plan can also be significantly enhanced around the following areas.

Strategic Structure

The 2019 Strategic Plan did a good job of discussing areas of concern and then sharing supporting challenges and R&D goals for each. This is a critical aspect for most-impactful National Science and Technology Council (NSTC) activities. The 2019 Plan's organization for doing so, however, led to some confusion. Primary aspects within the strategy (Framing, Defensive Elements, Priority Areas, and Critical Dependencies) were largely treated independently, when in reality there is often great interdependencies among them that (1) research sponsors would need to properly prioritize and scope their efforts, and (2) researchers would need to consider to optimize the return on investment from their activities.

MITRE therefore recommends a more strategically comprehensive approach to developing and organizing the 2024 Strategic Plan. It may help to use a strategic planning framework that is consistent with the Government Performance and Results Act.




Figure 1. Strategic Planning Framework with Values/Guiding Principles

Such a structured planning framework provides:

- A set of values and principles that guides all subsequent activities
- A universal and compelling vision for the future state of cybersecurity
- A series of goals that collectively enables the vision to be met
- Subordinate research objectives and strategies for each goal that are specific and time-bound and that both help drive activities to successfully meet goals and provide the Executive Office of the President (EOP) the ability to measure progress

Collaboration

The 2019 document took the unfortunately common approach of being very predominantly internal government focused when public-private collaboration is required on this topic both to determine research priorities and to optimally drive capability advancement and adoption.^{6,7} This has not been the case for a handful of most-successful prior NSTC activities,⁸ and MITRE strongly recommends (1) pointedly folding in nongovernmental entities into the strategy development and implementation process and (2) taking actions to help ensure that research outcomes are maximally leveraged (e.g., also supporting standards development, transition to manufacturing, considering supply chain ramifications).

Similar NSTC strategies for technologies identified in the 2019 Priority Areas have also been developed and/or updated, thus requiring updated analyses. While doing so, MITRE recommends more explicit linkage between this Strategic Plan and those other strategies rather than the generalized discussion in the current document:

- What are the similarities and differences when investigating each issue through different lenses (e.g., cybersecurity vs. AI research)? How do both communities understand each other and collaborate?
- How would advancements in cybersecurity research impact these other strategies, and vice versa?

Linkages between this R&D Strategic Plan and the overarching National Security Strategy⁹ should also be explicitly shown.

Supporting Small Entities

Cybersecurity is a nearly universal national concern, especially in the post-COVID remote connectivity environment, with incidents at one location creating cascading concerns at many others. Unfortunately, many (if not most) of the nation's cyber nodes are managed by small

⁶ C. Ford, et al. A "Horizon Strategy" Framework for Science and Technology Policy for the U.S. Innovation Economy and America's Competitive Success. 2021. MITRE, <https://www.mitre.org/sites/default/files/2021-11/prs-21-1440-horizon-strategy-framework-science-technology-policy.pdf>.

⁷ Mid-Decade Challenges to National Competitiveness. 2022. Special Competitive Studies Project, <https://www.scsip.ai/wp-content/uploads/2022/09/SCSP-Mid-Decade-Challenges-to-National-Competitiveness.pdf>.

⁸ D. Blackburn and M. Garriss. A National Science and Technology Council for the 21st Century. 2021. MITRE, <https://www.mitre.org/sites/default/files/2021-09/pr-21-2388-national-science-technology-council.pdf>.

⁹ National Cybersecurity Strategy. 2023. The White House, <https://www.whitehouse.gov/wp-content/uploads/2023/03/National-Cybersecurity-Strategy-2023.pdf>.

entities or individuals who do not have the knowledge or resources to ensure their cybersecurity protection. Research into finding optimal ways to identify and then guide these entities into taking necessary actions is recommended. Such research into “raising the floor” of the nation’s collective cybersecurity posture should not be forgotten.

Questions Posed in the RFI

2. What areas of research or topics of the 2019 Strategic Plan no longer need to be prioritized for federally funded basic and applied research?

None of the research topics included in the 2019 Strategic Plan have been “solved,” and they all warrant consideration to continue being included in the updated research strategy. The topics would benefit from more strategic development and organization, as we discussed in the preceding section under “Strategic Structure.”

3. What areas of research or topics of the 2019 Strategic Plan should continue to be a priority for federally funded research and require continued Federal R&D investments?

All topics within the 2019 Strategic Plan should continue, albeit with more specificity and targeted timelines. Doing so will both help drive activities and provide the EOP the ability to measure progress. Data-driven comments on several existing areas are included below to help develop this aspect of the strategy.

Use of Artificial Intelligence (AI) in Cyber Operations

Due to the growing number of threats to AI-enabled systems and the increasing use of AI in cybersecurity applications, the intersection of cyber and AI has become a critical area of research in recent years. From a national Cyber Science and Technology (S&T) Strategy point of view, several key research topics in the intersection of cyber and AI need further investment and investigation, including AI for cybersecurity, AI security, explainable AI, and privacy and data protection.

With the continued shortage of qualified cybersecurity professionals and the increasing volume and severity of attacks, there is a growing reliance on AI-driven cybersecurity defenses. A better understanding is needed of the implications of this increased reliance on AI for cybersecurity. For example, an AI defender must consider the system being defended and the mission it performs, and account for the uncertainties associated with sensing and action. It must also consider an attacker’s capabilities to adapt their behavior based on what they observe about the system and the defender. In AI security, the unique opportunities presented by leveraging advancements in AI have led to an “AI arms race” between the U.S. and its adversaries.

AI-enabled systems have novel threat surfaces and unique vulnerabilities far beyond those of traditional cyber systems, as shown in the MITRE ATLAS™ framework,¹⁰ which is modeled after and directly compatible with MITRE ATT&CK. These new AI threat surfaces are being

¹⁰ MITRE ATLAS™ (Adversarial Threat Landscape for Artificial-Intelligence Systems). 2023. MITRE, <https://atlas.mitre.org/>. Last accessed February 27, 2023.

increasingly attacked and exploited¹¹ in operational systems across industry and government, and many organizations that operate AI-enabled systems are unaware of new threats¹² that cannot be fully understood or mitigated using current cybersecurity techniques. As recommended by the National Security Commission on Artificial Intelligence's Final Report,¹³ organizations should stand up AI red teams to focus specifically on threats to systems incorporating AI capabilities.

With AI-enabled systems becoming more ubiquitous and complex, as in the recent adoption of ChatGPT, it is increasingly important to understand how those systems make decisions, the level of trust users have in those decisions, and measures of explainability of the systems. User adoption of AI systems will rely heavily on the trust users have in them. Concerns about privacy and data protection are paramount in the deployment of AI systems. AI systems rely on massive amounts of data, some of it private and sensitive, so developing methods to protect that data while still enabling the benefits of AI must remain a priority. The vulnerability of AI-enabled systems to open or insecure data sets requires additional scrutiny, particularly in recent developments of AI that rely heavily on semi- and self-supervised training methods. This concern is also true for many other autonomous systems, such as securing autonomous vehicles and the Vehicle-to-Infrastructure grid.

Smart Cities, IoT, and Critical Infrastructure

Several aspects of the Trusted Distributed Digital Infrastructure envisioned by the 2019 Federal Cybersecurity R&D Strategic Plan have yet to be realized and require further research investment.

The 2019 strategy speaks of pursuing design frameworks “that integrate safety, security, and privacy requirements, allowing system designers and developers to reason across all three domains concurrently.” It continues to be the case that those three domains are often treated as disjoint concerns in smart city product development, particularly when it comes to the handling of data streams from smart city devices. Other cyber physical domains are beginning to address the security and safety aspects in a more unified way. However, this is not yet the case with privacy. Developers need better tools for linking privacy risks to smart city device and system design choices, as well as improved means to incorporate privacy-enhancing technologies (PETs) into these systems.

Many domains are newly incorporating Internet of Things (IoT) devices to enable new use cases. These may present specialized cybersecurity needs. As one example, IoT consumer and healthcare technologies within the connected digital health ecosystem—including device manufacturers, health delivery organizations, service providers, and government—have proliferated to accommodate individuals' desire to age with independence and to enable data collection/integration between patient/consumer and provider. This ecosystem is vulnerable to cyber attacks, with the growing older population as a susceptible target. Age-friendly user design related to cybersecurity is an overlooked but increasing risk. The healthcare to home trend requires improved cybersecurity, privacy, and usability of connected devices. Research is needed

¹¹ Case Studies. 2023. MITRE ATLAS, <https://atlas.mitre.org/studies/>. Last accessed February 27, 2023.

¹² Market Guide for AI Trust, Risk and Security Management. 2023. Gartner, <https://www.gartner.com/en/documents/4022879>. Last accessed February 27, 2023.

¹³ The Final Report. 2021. The National Security Commission on Artificial Intelligence, <https://www.nscai.gov/2021-final-report/>. Last accessed February 27, 2023.

to incorporate assessment of cybersecurity vulnerabilities associated with age-related functional decline of users and to identify age-friendly security controls and mitigation.

Securing the hardware and software involved in smart city systems is another challenge area. Numerous solutions for security hardening and implementing trust and assurance have been developed over the course of many years of research. These advances, however, continue to find limited use in deployed systems. In our experience, the kinds of IoT and cyber physical systems involved in smart cities are the least likely to have incorporated such technology. In our work with device vendors, system integrators, and operators in this space, we have identified that many entities see the benefit of applying more advance security technologies to their products and systems but lack the resources or expertise to do so effectively. Research investment should be made on how to reduce the cost and complexity of applying and deploying existing hardware and software security techniques or, alternatively, developing easier-to-use methods with the goal of making them more accessible (economically and skills-wise) to developers and operators.

Cyber Resiliency

Notable progress has been made in cyber resiliency since the 2019 strategy. NIST guidance (Special Publication 800-160, Volume 2, Developing Cyber-Resilient Systems) was published in 2019 and revised in 2021. Cyber resiliency has been incorporated in federal and Department of Defense (DoD) systems planning and evaluation. Cyber resiliency is a unifying concept encompassing approaches such as segmentation, diversity, unpredictability, deception, and others. However, critical elements are still needed, such as a deep understanding of cost and effectiveness trade-offs, a robust range of available implemented resiliency mechanisms and resilient-by-default services, approaches to coordinate operation of resiliency across interconnected systems at scale, and development of use cases to support research in different domains. In addition, cyber resiliency designs are largely hand-crafted. More dynamic, automated capabilities are needed, both to adapt to changing situations and threats and to enable establishment and use of resiliency measures by staff with limited expertise.

The following are some important areas for research:

- **Address at the Foundational Level.** New laws such as the Bipartisan Infrastructure Law are making investments to improve sectors such as transportation in the critical concern areas of safety, climate, and workforce development. To achieve strategies that address these concerns, the sectors will be accelerating their adoption and integration of multiple technologies, with minimal focus on cybersecurity and/or privacy. Ensuring cybersecurity is addressed while keeping pace with such an acceleration of adoption is challenging. To mitigate this disparity, R&D efforts should focus on cyber resiliency strategies to ensure cybersecurity is addressed at the foundational level (e.g., strategies for moving to zero trust (ZT), cyber and privacy resiliency for smart cities). R&D efforts should also include developing methods for cyber resiliency at the systems integration level.
- **Orchestrated Cyber Resiliency in Evolving Systems.** Capabilities must be developed to coherently integrate and orchestrate cyber resiliency solutions (e.g., COTS, GOTS, emerging technology, innovative proactive defenses) into existing and emerging architectures, for scalability and to avoid creating gaps or new attack surfaces. As advanced cyber adversaries continue to evolve—and may be able to exploit

vulnerabilities installed in shared codebases or commonly used products—this orchestration needs to be ongoing; there is no “one and done” against the advanced persistent threat.

- **Dynamic Mission Resilience at Scale.** The challenge of ensuring secure and resilient operations exists at multiple scales and in a dynamic environment. Solutions exist for individual systems, services, and infrastructures, but often assume a static environment. Critical infrastructures and essential functions of large organizations depend not on an individual system but on systems of systems. An essential function's dependencies are themselves often dynamic, based on mission phases. The challenge is to orchestrate capabilities and practices dynamically across different constituent elements as they carry out mission functions. Next-generation capabilities will need to integrate mature resiliency capabilities (e.g., non-persistence, deception) with current commercial architectures (e.g., cloud services, zero trust architectures) and provide readily available, pre-packaged, threat-informed defensive courses of action to prevent cyber adversaries from achieving their goals. These capabilities will require sophisticated decision engines that manage and adapt resiliency mechanisms and that assist in orchestrating decisions made at different levels of enterprises and interconnected systems.
- **Innovative Proactive Defenses.** Cyber resiliency techniques and technologies that counter advanced adversaries need not be predicated on first detecting the adversary. New industry-provided technical capabilities (e.g., in ZT technologies and cloud services) provide opportunities for supporting proactive cyber resiliency, potentially in combination with enhanced ICAM (Identity, Credentialing, and Access Management), cryptographic obfuscation, and cross-domain filtering for high-trust segmentation. Attention should also be given to extending and adapting resiliency techniques to non-enterprise IT architectures (e.g., operational technology (OT), IoT, and industrial internet of things environments).
- **Transforming Governance and Analytic Foundations for Cyber Resiliency.** The transition to dynamic resiliency at scale involves shifts in governance, informed by threat-informed, risk-sensitive analysis methods. Guidance on effective and scalable implementation of advanced cybersecurity and cyber resiliency capabilities is needed to underlie definition, governance, and implementation of cybersecurity doctrine across enterprises involving on-premises, cloud, and multi-cloud IT environments. To enable and accelerate evolution from compliance-oriented toward threat-informed, risk-sensitive risk governance, and to maximize the usability of the guidelines and doctrine, tools (e.g., CREF Navigator™) and analytic methods for identifying cyber resiliency solutions and making risk management decisions must be refined and extended to apply to a broader range of system types and operational environments.
- **Cyber Resiliency Analysis for Small and Medium-Sized Organizations.** Small and medium-sized organizations depend on managed services (e.g., internet, cloud-based application services, backup/restore services) and pre-packaged application systems, often from niche providers. Methodologies are needed to enable small and medium-sized managed service providers and developers of commodity application systems to take advantage of cyber resiliency engineering and identify cyber resiliency capabilities to prioritize in their service offerings and applications. Such methodologies must be based on factors that have been found to be relevant in prior engineering and research efforts, and on repeatable, explicable, and extensible reasoning methods. Factors that can drive

the selection of cyber resiliency capabilities include prior investments and legacy technology (which can limit what can be implemented, but also offer functional capabilities for security, continuity of operations, and safety that can be leveraged for cyber resiliency), legal or regulatory requirements, architectural trends (e.g., ZT, convergence of IT and OT), physical constraints, and operational constraints (e.g., limited user expertise or attention). Research is needed into how to capture such factors, apply and adapt them to the setting of small and medium-sized organizations, and make the analysis extensible to different domains or sectors and operational environments.

- **Resilient-by-Default Cloud Services.** Default cloud configurations are set for quick deployment and are configured with only minimal Risk Management Framework or Federal Risk and Authorization Management Program (FedRAMP) security. Orchestration engines (e.g., Kubernetes, Terraform) help automate the infrastructure deployment, but they leave infrastructure configuration and application of mission-specific resilience and security to the client organization. Capabilities are needed to derive and efficiently combine cyber-resilient and optimal ZT-compliant software-defined networking/provisioning baselines for cloud environments, informed by operational need from criticality analyses.

Microelectronics Supply Chain

The U.S. economy, military, and commercial sector are highly dependent on advanced microelectronics (uE), but the globalization of the uE supply chain—from design through manufacturing—has introduced numerous threats to the security of these components in the form of hardware trojan horses (HTH), malicious modifications, and counterfeit parts. Counterfeit parts alone are estimated to cost semiconductor manufacturers more than \$7.5 billion annually in lost revenue. Additionally, the threat of HTH or other malicious modifications has not only cost the DoD billions of dollars for protection but also precluded many DoD programs from using state-of-the-art microelectronics due to a lack of trusted sources.

Privacy

PETs represent one set of a series of tools that can be used to protect data and minimize legal, privacy, and ethical risks.¹⁴ Modern PETs offer the potential to protect sensitive data while also helping government agencies achieve their mission goals. PETs should be considered in furtherance of recently issued Presidential Executive Orders and goals on Signals Intelligence data sharing, Securing and Protecting Access to Healthcare Services, Ensuring Responsible Development of Digital Assets, and collective efforts to use healthcare data to find a cure for cancer. Research regarding the different types of PETs and their effectiveness in different scenarios should be performed so that appropriate PETs are selected for adoption and implemented effectively. Also needed is further development of crypto-supported “zero-knowledge proofs for privacy” and “differential privacy.”

Privacy threats are currently not well understood, and privacy threat modeling is not actively included in risk management processes within many organizations. It is important to effectively assess privacy threats and use privacy threat information as input for PETs selection and privacy

¹⁴ Response of The MITRE Corporation to the OSTP RFI on Advancing Privacy-Enhancing Technologies. 2022. MITRE, <https://www.mitre.org/sites/default/files/2022-08/pr-21-01760-26-response-mitre-corporation-ostp-rfi-advancing-privacy-enhancing-technologies.pdf>.

operations overall. Privacy risk models need to expand beyond consequences (which most of them focus on) to characterize threats and vulnerabilities as well. MITRE is currently developing a Privacy Attack Taxonomy named PANOPTIC that will provide a standard structure for mapping privacy attacks that can be used to model privacy threats and facilitate privacy risk management, including PETs selection. Security risk modeling typically focuses on confidentiality-based threats to information about individuals (e.g., data breaches). However, the Privacy Attack Taxonomy will enable identification of threats beyond those typically addressed in security risk modeling (e.g., threats related to consent; notice; and inappropriate use, sharing, or retention of information about individuals). This expansion in focus will also enable consideration of a broader set of PETs for potential implementation. Further research is needed in the area of privacy threat modeling and its implementations for privacy risk management.

AI ethics and privacy are critical considerations when developing and implementing artificial intelligence systems. As AI becomes increasingly integrated into daily life, it is essential to ensure that these systems are designed and used in ways that are ethical and transparent, and that respect people's privacy. Ethical considerations might include ensuring that AI algorithms do not discriminate against certain groups of people, protecting people's autonomy and dignity, and ensuring that AI systems are used for the common good. Privacy considerations might include protecting people's personal data, ensuring that data is collected and used only with individuals' consent, and ensuring that data is stored and transmitted securely. As AI continues to advance, research is needed to identify how best to address these critical ethical and privacy concerns to ensure that AI is used in a way that benefits society as a whole.

The IoT raises a host of privacy-related questions in need of more systematic exploration. These are difficult issues because they are cross-cutting. The architectural progression from lightweight sensors to sophisticated analytical systems complicates matters due to the significant shift in properties at almost every stage of the pipeline. The nature and implications of such shifts in terms of privacy have not been well analyzed, much less tackled. Emergent properties are likely at these points, including the social effects of implicitly infrastructural surveillance. Existing privacy risk models—even the more sophisticated ones—struggle to effectively capture these kinds of concerns. Privacy-enhancing technologies cannot be relied on to ride to the rescue as their use cases tend to be narrowly focused, lacking any calculus regarding composability. Approaches to data and AI ethics, meanwhile, tend to emphasize technical correctives and procedural remedies while side-stepping more fundamental questions regarding power asymmetries and the effects of distributed yet interacting decisions, among other issues. Holistically addressing privacy in IoT environments requires more formal and rigorous investigations into all of these topics, but from distinct vantage points.

4. What objectives not included in the 2019 Strategic Plan should be strategic priorities for federally funded R&D in cybersecurity? Discuss the challenges, desired capabilities and outcomes, and objectives that should guide research to achieve the desired capabilities, and why those capabilities and outcomes should be strategic priorities for federally funded R&D.

Defensive Capabilities for Under-Resourced Entities

Ransomware attacks on mission/business-critical operations have increased over the past four years, with many organizations unable to recover from them. While small and medium-sized businesses (SMBs) tend to believe they are too small to be targeted by ransomware (or other attacks), evidence does not support this belief. Furthermore, approximately one-third of SMBs in the U.S. reported they had to close following a ransomware attack. Those that did not shutter suffered significant loss of revenue.¹⁵ This is a disturbing trend, especially given that small businesses generate 44 percent of U.S. economic activity.¹⁶ Research that develops knowledge and tools to help these entities should be considered.

Establishing Partnerships with the Explicit Purpose of Enabling Assurance Evaluations of Large ML Models

Large machine learning (ML) models are increasingly applied to solve problems such as drug discovery¹⁷ and protein folding,¹⁸ with broad implications for humanity. However, the negative implications of generative models are not yet well understood. For example, models could be used for generating toxic molecules.¹⁹ Similarly, large models that enable the generation of text^{20,21} and source code^{22,23} at scale could be used to generate adverse content to manipulate people or create malware, respectively.

Tools to Understand the Pedigree of Software and Greater Visibility of Its Supply Chain

Vulnerabilities in third-party software are the root cause of one of the most expensive cyber incidents on average, that organizations endure.²⁴ Research to understand and support software supply chain security should be considered including the need for capabilities such as software bills of materials (SBOMs) and how to adopt/implement their usage as well as methods and mechanisms an organization can use to determine whether appropriate choices were made for securing the software during the creation process. This includes information about the compilation and formulation options used in transforming the source components and parts into the resultant software.

¹⁵ Ransomware: The True Cost to Business. 2022. Cybereason, https://www.cybereason.com/hubfs/Ransomware_True_Cost_e-book_NewBrand.pdf.

¹⁶ Small Businesses Generate 44 Percent Of U.S. Economic Activity. 2019. U.S. Small Business Administration, <https://advocacy.sba.gov/2019/01/30/small-businesses-generate-44-percent-of-u-s-economic-activity/>. Last accessed March 1, 2023.

¹⁷ F. Urbina, et al. Dual Use of Artificial-Intelligence-Powered Drug Discovery. 2022. Nature Machine Intelligence, <https://www.nature.com/articles/s42256-022-00465-9>. Last accessed February 27, 2023.

¹⁸ J. Jumper, et al. High Accuracy Protein Structure Prediction Using Deep Learning. 2020. Fourteenth Critical Assessment of Techniques for Protein Structure Prediction (Abstract Book).

¹⁹ F. Urbina, et al.

²⁰ N. Stiennon, et al. Learning to Summarize from Human Feedback. 2020. 34th Conference on Neural Information Processing Systems, <https://proceedings.neurips.cc/paper/2020/file/1f89885d556929e98d3ef9b86448f951-Paper.pdf>.

²¹ J. Devlin, et al. BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding. 2018. ArXiv, <https://arxiv.org/pdf/1810.04805.pdf&usq=ALkJrhzhxICL6yTh2BRmH9atgvKFxHsxQ>.

²² Y. Li. Competition-Level Code Generation with AlphaCode. 2022. Science, <https://www.science.org/doi/10.1126/science.abq1158>. Last accessed February 27, 2023.

²³ OpenAI Codex. 2021. OpenAI, <https://openai.com/blog/openai-codex/>. Last accessed February 27, 2023.

²⁴ Cost of a Data Breach 2022. 2022. IBM Security, <https://www.ibm.com/reports/data-breach>. Last accessed February 27, 2023.

Executive Order 14028 has motivated much progress toward requiring attestations from software suppliers regarding which of the items in NIST's SP800-218, Secure Software Development Framework, an organization applied when developing its software. However, there are no standards for how the software developers can be captured and conveyed in a manner that others can understand, machines can act on, be linked to the SBOM for the software, and that consumers of the software will have a reason to trust the information and its source.

There are promising approaches in the Linux Foundation's in-toto project as well as the Internet Engineering Task Force Supply Chain Integrity, Transparency and Trust working group. However, many open questions remain requiring critical thinking and evaluation of these proposals in order for the U.S. government and industry to rely on them for their future software and dynamic supply chains. Accelerating the convergence and maturity of these complementary efforts will address a massive capability gap that sorely needs to be addressed to achieve an overall articulated framework that can provide assurance of software and can convey the appropriate information tailored to the situation the software will be used in. An evidence-based approach that allows for self- and third-party verification in a highly automated manner would provide an evidence chain to demonstrate and measure trust of the pedigree and provenance of the software supply chain.

Increasing Hardware Assurance

The foundation of trust in computer systems is built on trust in hardware, and trust in hardware is in turn based on trust in hardware components that are of critical importance to the properties of the computer system. An important goal in advancing the assurance of computer systems is to constantly look for architectures that can be trusted based on fewer or simpler hardware or software components.

Also needed are hardware assurance primitives that are easy and cost-effective to integrate into industrial control systems, vehicles, and IoT devices. Poorly understood cost-benefit continues to be a barrier to adoption.

Cyber Deception and Adversary Engagement

Cyber deception, when used as an integral element of the defensive cyber framework, provides benefits across the goals of deter (diminish value of spoils), protect, and detect by providing high-quality indicators of compromise and techniques used to exploit and in current manifestations provide information to defenders to initiate response options. As AI/ML is increasingly used in attacking systems, deception can be utilized to counter and respond to increasing sophistication and speed of attacks. Needed capabilities include:

- Dynamic adaption of a deception environment to vary the perceived attack surface
- Providing deeper adversary attack opportunities while yielding no valuable data or exposure of protected systems
- Utilization of hyper-instrumentation from the deception environment to create widely available defensive signatures in real time to accelerate both detection to the locally defended system and distribution to standard defensive systems that may not deploy cyber deception
- Generation of response options to dynamically change the operational configuration of the defended system while maintaining the adversary's perception of the environment

- Real-time generation of defensive options that could be applied locally to thwart or render an attack ineffective
- AI-driven adversary behavior to classify and evaluate cyber deception techniques to advance testing and evaluation of deceptive capabilities
- Generation of defensive options that could be distributed to service providers to change the engagement or redirect an attack

Individually, each capability advances the objectives of the defensive framework pillars; combined, they provide the ability to change the defensive posture from one of passive response to active engagement to reduce the effectiveness of an attack.

5. What other scientific, technological, economic, legal, or societal changes and developments occurring now or in the foreseeable future have the potential to significantly disrupt our abilities to secure the digital ecosystem and make it resilient? Discuss what federally funded R&D could improve the understanding of such developments and improve the capabilities needed to mitigate against such disruptions.

Preparing for the Post-Quantum Crypto Migration

Over the next 5–10 years, it will become increasingly important to move beyond today's factoring or discrete logarithm-based asymmetric cryptography toward alternative methods collectively known as "post-quantum cryptography." Migrations of cryptographic algorithms are notoriously slow and difficult. Among topics particularly relevant to post-quantum migration include protocol-oriented and standards-oriented study designed to anticipate and remove obstacles to adoption, as well as to minimize the impacts of increased resource demands from post-quantum algorithm use. Additional challenges include:

- Developing and improving post-quantum versions of alternative asymmetric crypto primitives such as identity-based encryption, attribute-based encryption, and so on
- Maturing post-quantum fully homomorphic encryption and developing a standard framework for expressing homomorphic transformations to enable applications

Embedded Security (Side-Channel, Fault Injection, etc.)

Embedded systems security is essential to the civilian, defense, and intelligence communities in their efforts to secure embedded systems. Embedded systems are the backbone of all modern infrastructure, sensing, navigation, communication, and weapons system capabilities. These capabilities have been applied to secure infrastructure and end-user equipment for a broad set of areas, including mobile, medical, transportation, and navigation. New R&D is necessary to better understand emerging threats and develop defensive technologies and tools to protect these systems and combat supply chain threats.

Zero Trust

ZT principles are being adopted broadly to reduce exposure and limit the ability for attacks to move freely within enterprises. Research is needed to advance the sophistication and effectiveness of architectures incorporating ZT principles, including:

- Application of ML and AI to ZT capabilities such as automation, threat detection, vulnerability management, access decision, risk evaluation, and detection. Additional research will be required in this area to stay ahead of the expected use of AI, ML, and automation by adversaries.
- Approaches for interoperability between ZT Policy Enforcement Points (PEPs), orchestration platforms, services, and the like, which could include needed standards for information exchange between platforms and interpretation of telemetry.
- ZT approaches for architectures expanding beyond information technology to incorporate OT, 5G and other advanced networking capabilities, and Web 3.0/Metaverse.
- Dynamic policy enforcement and enhancing interoperability between ZT products and architectures, including the ability to communicate the point-in-time risk posture of subjects and resources between ZT components.
- Ways to fully incorporate access-relevant information and PEPs available from applications into end-to-end ZT risk evaluation and enforcement to enable more data-centric security.

6. What further advancements to cybersecurity education and workforce development, at all levels of education, should be considered to prepare students, faculty, and the workforce in the next decade for emerging cybersecurity challenges, such as the implications of artificial intelligence, quantum computing, and the Internet of Things on cybersecurity?

MITRE has previously responded to an Office of the National Cyber Director (ONCD) RFI that was wholly focused on this specific question.²⁵ Our overarching recommendation within that response was that maximizing growth within the cyber career path requires targeted and cohesive development throughout its pipeline while simultaneously recognizing that staff may onboard and leave at various stages. MITRE conceptually views this pipeline as depicted in Figure 2, which could easily be genericized for application at the national level and complement the suite of cyber workforce-related tools and approaches already offered by the Office of Personnel Management and other federal agencies.²⁶ It could also serve as a model for the private sector, though some adaptation will be required for smaller, less-technical organizations.

²⁵ MITRE's Response to the ONCD RFI on a National Cyber Workforce Strategy. 2022. MITRE, <https://www.mitre.org/sites/default/files/2022-11/pr-22-01891-09-mitres-response-oncd-rfi-national-cyber-workforce-strategy.pdf>.

²⁶ For example, see Workforce Planning for the Cybersecurity Workforce. 2023. Office of Personnel Management, <https://www.opm.gov/policy-data-oversight/human-capital-management/cybersecurity/>. Last accessed March 2, 2023.

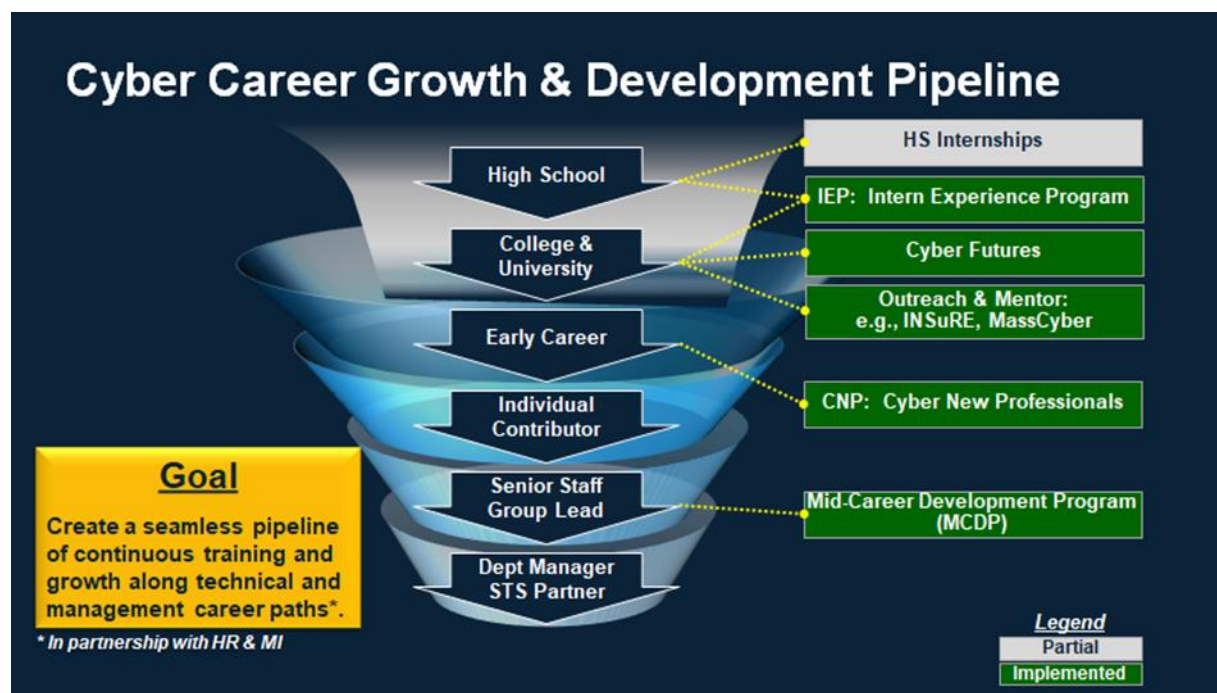


Figure 2. Cyber Career Growth and Development Pipeline

We have developed learning programs at, and specific to, multiple stages within this pipeline, which are discussed in the response to the Office of the National Cyber Director RFI..

MITRE has also developed a Cyber Workforce Development Framework to help nations build a cyber workforce strategy. This framework is being applied on behalf of the U.S. State Department in several partner nations around the world, and by the Department of Homeland Security's Cyber Development, Education, and Training program in response to Solarium Commission recommendations.²⁷ Findings indicate that governments are usually best positioned to integrate the cyber ecosystem of industry, academia, cyber professionals, commercial training programs, and national security needs. They can convene leaders, develop standards, set long-range goals, incentivize progress and cooperation, and eliminate barriers. But government is not enough. The main finding of the framework survey is that public-private partnerships are key. Together, governments and private organizations can develop standards for training, job descriptions, and career pathways; develop policies that facilitate workforce development; and identify regional barriers for cyber career seekers. This framework can be used to collaboratively identify educational and workforce development needs for a wide range of entities:

- Nations transitioning to a digital economy or adjusting incentives and pipelines to increase investment in and access to cyber professionals
- City, state, or regional planning groups focused on increasing high-tech employment and an associated ecosystem
- Industry and academia seeking to grow a local talent pool

²⁷ Construction of the framework began with a broad survey of tech workforce development approaches across nations of various sizes and economies, development non-governmental organizations, and other subject matter experts. They identified commonalities, needs, and best practices in several categories and then synthesized the information into the framework, focused on key areas and approaches.

- Government agencies at all levels developing policy and/or legislation to incentivize cyber talent development and retention in key functional areas

Workforce Development through Gamification

Gamification of hands-on project-based learning activities is a proven technique for significantly improving learning outcomes and making complex subjects more accessible to a wider range of students. Since 2015, MITRE has been perfecting this technique as applied to its Embedded Capture-the-Flag (eCTF) event—a nationwide competition for high schools and colleges that provides an unmatched learning experience for embedded systems security (i.e., Secure Edge/IoT computing).

Embedded systems (like those found in smartphones, modern automobiles, and many military systems that are critical for national security) are significantly different from other computing systems. As such, these systems face unique security threats that require different skillsets than those needed for addressing traditional cybersecurity. MITRE noticed an educational gap in U.S. schools and universities in curricula for embedded systems security—and this gap largely still persists today.

To reduce this gap, MITRE developed the eCTF competition in 2015, and has successfully used it to raise awareness of this important field of study and to spur an interest among the nation's student population. The competition design was based on existing “capture-the-flag” competitions that have become popular for teaching traditional cybersecurity, but with three significant differences that make the eCTF unique:

1. A focus on embedded systems. Teams implement designs and conduct attacks on software running on real physical hardware.
2. Attack-and-defend. Teams design and implement their own solutions to the challenge and then develop attacks against the designs of other teams.
3. Extended time. The competition runs for three months over the spring semester from mid-January through mid-April to allow time for teams to design, implement, and attack.

Since the first year with only four universities, the competition has grown to 80 schools and over 500 students ranging in level from high school to Ph.D. Participant surveys indicate that the competition has greatly influenced students toward pursuing careers in embedded security and related fields.

The eCTF has also proved to be a powerful tool in diversity, equity, and inclusion. Through targeted outreach, the eCTF has been able to engage with three Historically Black Colleges and Universities and eight other Minority Serving Institutions, engaging communities historically underrepresented in embedded security and adjacent fields. After their inaugural participation in the eCTF, professors at Morgan State University published a paper about their success at leveraging the eCTF to attract minority students to their programs in secure embedded systems.²⁸

²⁸ M. Kornegay, et al. Engaging Underrepresented Students in Cybersecurity Using Capture-the-Flag(CTF) Competitions (Experience). 2021. ASEE, <https://peer.asee.org/engaging-underrepresented-students-in-cybersecurity-using-capture-the-flag-ctf-competitions-experience>. Last accessed February 27, 2023.

Appendix A. Overview of MITRE's Cybersecurity Activities

Among MITRE's most renowned capabilities, cybersecurity is a core competency demonstrated by decades of seminal innovations and contributions to advancing the field. Throughout our 50-plus year history in cybersecurity, we have earned recognition for game-changing advances gained through collaborative processes and methods, starting with major contributions to the first government series, the Orange Book and Rainbow Series. Most innovations are widely adopted and used today, from the original designs of cross-domain architectures using MITRE's Bell LaPadula model to CVE®, the taxonomy used to characterize vulnerabilities, to more recent innovations such as ATT&CK® and CALDERA. We work across government, industry, and academic partners to develop, identify, and adopt new concepts and apply threat-informed engineering to build and defend resilient environments. In addition to MITRE's own innovations, we draw from other organizations' best-of-breed solutions and capabilities to bring them to our sponsors.

MITRE has broad and deep knowledge of secure and resilient architectures, cyber technologies, and cybersecurity operations that enable us to knowledgeably approach a comprehensive range of cyber-related challenges. In addition, MITRE has deep strength in the areas of cryptography, privacy, and cyber supply chain security, and continues to transform vulnerability management and threat intelligence.

MITRE works across the national security and civil sectors and industry to advance secure architectures and defensive cyber operations, develop innovative cybersecurity solutions, and analyze the cybersecurity implications of new and emerging technologies and applications.

MITRE leverages these diverse areas of cybersecurity expertise in a multidisciplinary perspective to assess emerging events such as SolarWinds, an attack that combined multiple direct and supply chain exploit strategies.

Innovation Highlights

MITRE balances classic cyber defense approaches and innovation with a strong emphasis on leveraging cyber threat intelligence to respond and adapt quickly to cyber attacks. To accomplish this, we form partnerships that promote sharing cyber threat information and effective tools. Our strategy thrives on a foundation of unrelenting innovation and operational experimentation. MITRE's cyber expertise includes wireless and wired network security, mobile device security, threat intelligence and hunting, cybersecurity assessments, Internet of Things security, medical device security, and adversary emulation, among other areas. Examples of MITRE's cyber security innovations are listed below.

Originator of ATT&CK® – Adopted by Industry and Government

MITRE is the creator and maintainer of the game-changing [ATT&CK framework](#), which codifies the cyber attack tactics, techniques, and procedures known to be used by sophisticated adversaries. ATT&CK is populated by validated reports of adversary attack techniques contributed by a worldwide community of cybersecurity experts, creating synergy that favors the defender. ATT&CK quickly became the international nexus of public information about techniques revealed by adversary attacks and exploits. It has been used in many organizations as a basis to develop new detection capabilities and to analyze the coverage of their cybersecurity

sensors and countermeasures. As an example, the Department of Homeland Security (DHS) Cybersecurity and Infrastructure Security Agency (CISA) recently released a guide, titled [Best Practices for Mapping to MITRE ATT&CK®](#), in an effort to encourage common language across organizations for threat identification and analysis. ATT&CK covers attack techniques against enterprise information technology, mobile technology, and industrial control systems, and its scope continues to broaden.

Originator of CVE®: The Taxonomy and Registry for Cybersecurity Vulnerabilities

MITRE originated CVE, the universally relied-upon [Common Vulnerabilities and Exposures registry](#), and maintains it for CISA. Working with industry, MITRE tracks, validates, and publishes reported vulnerabilities. This MITRE innovation played an essential role in creating a cohesive cybersecurity community by enabling unambiguous communication among cybersecurity experts, vendors, and user organizations about what vulnerabilities exist, in what software they are found, and what cybersecurity mitigations can counter them. MITRE also created and maintains CWE™, the [Common Weakness Enumeration](#), which catalogs the types of weaknesses in code that can lead to security vulnerabilities, with active input from the cybersecurity community. This allows recognition of weaknesses that are present in specific software components and creation of acquisition requirements specifying what must be eliminated. CWE publishes an annual list of the top 25 most dangerous software weaknesses as a way of drawing public attention to the most urgent issues to fix. CWE also recently, in partnership with Intel Corporation, expanded its scope to include hardware weaknesses.

Operator of NCF, the Nation's Only Cyber-Focused FFRDC

MITRE operates the nation's only cyber-focused FFRDC, the [National Cybersecurity FFRDC](#) (NCF), in support of the National Institute of Standards and Technology ([NIST](#)) [Cybersecurity Center of Excellence](#) (NCCoE). NCF's mission is to enable strong, practical, and resilient cybersecurity for all Americans. Sponsored by NCCoE, NCF advances the state of cybersecurity practice across industry, infrastructure owners/operators, commercial solution providers, government, and academia. With multiple laboratories to create exemplary integrated cyber solutions in topics ranging from Zero Trust Architecture to identity and access management, to medical devices, to commercial space systems, NCF develops practice guides and explores a range of pressing cybersecurity needs.

Leader in Cyber Resiliency

MITRE has been at the forefront of establishing the field of cyber resiliency for the past decade. MITRE built a community across government, industry, and academia, in part by creating seminal artifacts and by establishing annual workshops. The Cyber Resiliency Engineering Framework (CREF) developed by MITRE is the comprehensive technical framework providing guidance for implementation of cyber resiliency promulgated in NIST Special Publication (SP) 800-160 Vol. 2 (Rev 1): [Developing Cyber-Resilient Systems](#). The CREF provides the keys to making systems across the government and more broadly in the private sector more resilient to cyber compromises and attacks, allowing them to continue accomplishing their missions without disruption in the face of sophisticated cyber attacks.

Recently, we held ResilienCyCon, which brought together experts from government, industry, and academia to discuss how cyber resiliency has been adopted in different domains and fostered

synergy across the community. Industry, government, and academia leaders have adopted MITRE's Cyber Resiliency approach, as evidenced by speakers including the Google Cloud CISO; Amtrak CISO; Options Clearing Corporation Chief Security Officer; Ariam VP for Cybersecurity and Incident Response; Head of Cyber Resiliency at Standard Chartered Bank (in Poland); and senior government officials from OSD, Army, and NIST. Internationally, as part of the [IEEE Conference on Cyber Security and Resilience](#) (CSR), MITRE has partnered with PNNL for the past six years to co-lead the [CSR Workshop on Cyber Resilience and Economics](#).

The [CREF Navigator](#)TM, which makes the CREF more usable and accessible, is a major contribution in helping to bring its benefits to more organizations. CREF Navigator is a tool for effective visual browsing, distilling complex concepts and relationships into understandable tailored views to enable architectural and engineering discussions and analysis (see [MITRE Launches Cyber Resiliency Engineering Framework Navigator](#)). Ron Ross, NIST Fellow, regularly posts about the importance of the CREF and the CREF Navigator to enhance the resilience of the nation to cyber attacks.

Thought Leader in Health Cybersecurity

MITRE has provided cybersecurity support to the Food and Drug Administration's (FDA's) Center for Devices and Radiological Health since 2014. To help FDA understand the different stakeholder perspectives in coordinated vulnerability disclosure, MITRE conducted a stakeholder study. MITRE co-authored with FDA a journal article summarizing [The Evolving State of Medical Device Cybersecurity](#) in AAMI Biomedical Instrumentation & Technology.

MITRE developed the [Playbook for Threat Modeling Medical Devices](#) based on a series of threat modeling bootcamps sponsored by FDA and conducted by MITRE, the Medical Device Innovation Consortium, to encourage the adoption of threat modeling in the design and development of medical devices. Threat modeling is an element of FDA's new draft premarket guidance, and the playbook offers examples that medical device manufacturers can use to develop threat modeling training and aid in the adoption of threat modeling best practices.

MITRE is also working with hospitals to prepare for and respond to ransomware and other cyber attacks, since they are highly targeted. Hospitals develop and exercise emergency response plans for all kinds of emergencies, but we are helping them include cyber incidents, which typically involve longer downtimes and more widespread disruptions compared with other incidents. MITRE recently published the [Medical Device Cybersecurity Regional Incident Preparedness and Response Playbook](#) to capture some best practices. To help hospitals and others deal with ransomware threats, we created the Ransomware Resource Center (<https://healthcyber.mitre.org>) and tailored a lightweight assessment methodology Cyber Operations Rapid Assessment (CORA) for healthcare delivery systems ([Threat-Informed Cybersecurity Operations for Healthcare Delivery Organizations](#)).

MITRE has provided technical analysis to support FDA's regulatory role and has developed tools and documents to help medical device manufacturers implement FDA's guidance, which represents its current interpretation of the regulations. MITRE developed a rubric for applying the Common Vulnerability Scoring System to medical devices that considered the clinical context, which FDA qualified as a Medical Device Development Tool, and helps medical device manufacturers assess the exploitability of vulnerabilities.

Leader in Cyber Supply Chain Security and System of Trust

MITRE has been influential in developing approaches to improve the security of the software supply chain against being used to infiltrate cybersecurity attacks and malware into our systems. MITRE was one of the original contributors to the concept of a software bill of materials (SBOM) to capture and track the supply chains inherent in widely used software products that organizations use. SBOMs are now mandated in the May 2021 Cybersecurity Executive Order 14028, and MITRE is assisting early adopters in developing the first SBOMs. MITRE has also developed the System of TrustTM framework that defines attributes of suppliers, services, and products that should be scrutinized to assess supply chain risk, as well as mitigations that can be instituted to reduce risk.

Creator of CALDERATM – Adversary Emulation to Automate Assessments

Seeing a need to provide tools that test by using the same TTPs adversaries use, MITRE developed [CALDERA](#), a cybersecurity framework that empowers cyber practitioners to develop and use automated security assessments on a scalable, automated adversary emulation platform. The platform is used by red teamers to develop tests, as trainers for red teams and defenders, for testing defensive tools, and for assessing networks through “real” simulated (controlled) adversary attacks. The platform is being extended to support operational technology.

Originator of EngageTM – Advancing Tools for Adversary Engagement

MITRE has developed a collection of resources centered on [MITRE Engage](#), a matrix designed to give decision makers and defenders the tools and common language they need to plan and analyze adversary engagement. Adversary engagement is an iterative, goal-driven process that leverages cyber denial and deception in unison to drive strategic planning. Unlike other defensive technologies, such as antivirus, adversary engagement technologies cannot be considered “fire and forget” solutions. Rather, an organization must think critically about what their defensive goals are and how denial, deception, and adversary engagement can be used to drive progress toward these goals.

Thought Leader in Defensive Cyber Operations

MITRE combines cutting-edge research with direct experience and threat intelligence from live environments to establish and evolve operations centers. We engineer solutions that work in real operations, identify where to invest scarce resources to optimize defense, and assess the ability to defend against real adversaries. Last year, MITRE published a book—[11 Strategies of a World-Class Cybersecurity Operations Center](#)—based on lessons learned across multiple government sponsors.

Creator of SAF – Automating Security Assessments and Baselines

The MITRE [Security Automation Framework \(SAF\)](#) is a set of tools designed to assist developers and security professionals in assessing systems against standards or custom baselines, apply automated mitigations, and ultimately harden systems. SAF provides open-source tools that help build security directly into software. This includes tools for developing security content based on specific security requirements as well as visualizing the security testing of any tool and trending over time. SAF is used at MITRE and in several software pipelines including the

Department of Defense, DHS, the Intelligence Community, and others. Several commercial companies have adopted SAF as well.

Additional Cyber Capability Highlights

MITRE operates special-purpose laboratories where we assess, engineer, and integrate technologies and processes to fill gaps and create affordable solutions to customer problems. Our laboratory capabilities include tools and techniques for evaluating and enhancing enterprise-class cyber defense products, applied mission resilience techniques, cyber-physical security technologies, cloud technologies, and cross-domain solutions.

Zero Trust and Cloud Architectures

MITRE is working across the federal government to evolve to Zero Trust Architecture principles and to secure sensitive data and applications in public and private clouds. We develop secure and resilient solutions to integrate cloud services with enterprise on-premises systems and business services. This includes taking an end-to-end perspective to ensure protection of data-at-rest, data-in-transit, data-in-use, and ransomware mitigation, and assessing appropriate security controls as commercial Cloud Service Providers (e.g., AWS, Azure) act as “scalars” to become 5G system and service providers. Expertise ranges from developing best practices in cloud security, design, migration, and operations to evaluating Cloud Service Providers, to implementing cyber analytic cloud environments.

MITRE has developed the CAVEAT (Cloud Adversarial Vectors, Exploits, and Threats) framework to capture cloud-specific adversary behaviors and potential exploits. CAVEAT will be evolved and maintained by the Cloud Security Alliance in collaboration with MITRE.

Cybersecurity Strategy, Policy, and Governance

MITRE has worked with many government organizations to develop and evolve their cybersecurity strategies and implementation roadmaps with focuses ranging from enterprise modernization to cybersecurity science and technology. In emerging nations, MITRE has worked with the State Department to develop national strategies to advance cybersecurity in their organizations, practices, workforces, and populations.

Cybersecurity Guidance

Historically, MITRE has been a key contributor to cybersecurity guidance since the earliest guidance created by the National Security Agency and NIST. MITRE has provided input to or co-authored guidance including NIST's special publications on security and privacy controls (NIST SP 800-53), system security engineering, and cyber resiliency (NIST SP 800-160 Vols. 1 and 2). In addition, MITRE has developed security profiles for applying controls to specific domains, assists organizations in implementing the guidance in their enterprises, and has supported government assessment of vendor products against guidance in programs such as FedRAMP and National Information Assurance Partnership. MITRE also performs technical analysis on implications of cyber technology-related decisions or policies for organizations developing guidance or regulations.

5G Cybersecurity

While 5G seeks to be more secure than prior mobile networks, its virtualization/cloud usage and service-oriented architecture still present a significant attack surface that must be carefully considered. MITRE has studied the cybersecurity aspects of 5G and has developed a threat-based framework for 5G security and resiliency, FiGHT™ (Five-G Hierarchy of Threats). FiGHT covers 5G components, critical assets, threat vectors, and threat actors. It is intended to provide a comprehensive basis for understanding risks from known and emerging threats, in standards or in specific architectures, to be able to identify mitigations and reduce the attack surface. FiGHT, which has just been released publicly, enables the 5G ecosystem to build, configure, and deploy secure and resilient 5G systems for specific use cases and architectures, giving carriers, service providers, and enterprises the tools to quantify risks, share threat intelligence, and plan for cyber investments. MITRE is also engaged in analyzing and improving the cybersecurity of applications built on increasingly capable mobile communications and special-purpose devices, such as Internet of Things (IoT) and smart cities.

Other cyber technical capability areas include:

- **Cyber Assessments:** We tailor and apply a full range of vulnerability, architectural, and adversarial assessment methods to unique sponsor needs and technologies across the system life cycle.
- **Privacy:** We address sponsor privacy concerns systematically by developing privacy programs and strategies, ensuring compliance, and providing training. MITRE has written two books on privacy that are widely used across the federal government.
- **Mobile:** We develop technologies and approaches to securely integrate mobile devices and apps into enterprise and mission environments, based on threat assessments and adversary emulation.
- **Cyber Physical Security and Internet of Things:** We develop approaches to allow safe and secure use at scale of cyber-physical and Internet of Things systems that operate with varying degrees of autonomy and have physical consequences.
- **Cross-Domain Solutions:** We help sponsors securely access and transfer data across security domains through effective combinations of technology, architecture, and policy.
- **Identity and Access Management:** We develop scalable, interoperable COTS-based approaches to securely manage identities and perform authentication and access control throughout an enterprise.
- **Trust and Assurance:** We advance the state of the art through R&D in software assurance, embedded systems analysis, cryptographic protocols, and trusted computing.
- **Cyber Workforce:** We create and implement a range of innovative approaches to developing the cyber workforce, including an Intern Experience Program, Cyber New Professionals (for early career employees), neurodiversity, mid-career development, and others.
- **Cyber Threat Intelligence and Sharing:** We provide actionable knowledge of and insight into adversaries and their malicious behaviors in order to inform traditional and non-traditional cybersecurity defensive and offensive missions by providing better visibility, reducing harm, and enabling better security decision making via an iterative, repeatable process.

- **Cyber Data Analytics and Malware:** We apply state-of-the-art data analytics to cybersecurity problems, including threat detection and in-depth analysis.
- **Cyber Effects and Reverse Engineering:** We enable offensive cyber operations with software and hardware techniques and solutions. Our evaluations, rapid prototypes, and tools provide essential capabilities designed for real-world missions.
- **Cyber for Critical Infrastructure Protection:** We provide a range of cybersecurity capabilities for the detection, protection, and defense of critical infrastructure operational and information technologies.

MITRE has performed 650+ cyber assessments across 26+ agencies within the Federal Civilian Executive Branch, Department of Defense, and Intelligence Community, and has experience assessing systems that span on-premises, remote, cloud-based, and hybrid architectures; mobile and IoT; and specialty medical devices.

In addition to conducting cyber assessments across the federal government, MITRE has performed significant work with leading cybersecurity guidance and policy organizations including NIST, CISA, and the Office of the National Cybersecurity Director (ONCD) in the Executive Office of the President. This includes working with CISA on software supply chain planning guidance and assisting ONCD with 2024 cybersecurity budget guidance. We are also serving as a trusted adviser to the Office of Management and Budget's (OMB's) Chief Information Security Officer, providing input on recent policies, including zero trust and software supply chain risk management. Congress asked MITRE for input into the Federal Information Security Modernization Act (FISMA), and last year MITRE testified on changes needed to the Federal Information Technology Acquisition Reform Act scorecard, offering cybersecurity recommendations aligned with the administration's direction in its Cybersecurity Executive Order and Zero Trust Architecture guidance. Our recommendations helped inform the cybersecurity metrics OMB has recently issued on performance.gov and for FISMA reporting, which will be the foundation for future scorecard grades.