CLOUD & INFRASTRUCTURE SUMMIT

JUNE 25, 2019 | MARRIOTT METRO CENTER | WASHINGTON, D.C.

On behalf of the Advanced Technology Academic Research Center, I am proud to announce the release of a White Paper documenting the MITRE-ATARC Cloud & Infrastructure Collaboration Symposium held on June 25, 2019 in Washington, D.C. in conjunction with the ATARC Federal Cloud & Infrastructure Summit.

I would like to take this opportunity to recognize the following session leads for their contributions:

MITRE Chair: Justin Brunelle, Principal Researcher, MITRE

Challenge Area 1: Cloud Smart & Impact on Cloud Adoption
Russell Pavlicek, Cloud Solutions Architect, Public Sector, RedHat
Stephen Kovac, VP of Global Government and Compliance, ZScaler
Katy Warren, Principal Engineer, MITRE

Challenge Area 2: Challenges to Cloud Enabling Security
Michael Fugate, Technical Manager / Solutions Architect, Pyramid Systems
Dan Tucker, Vice President, Booz Allen Hamilton
Mari Spina, Principal Cybersecurity Engineer, MITRE

Challenge Area 3: Cloud-enabled Rapid Development and DevSecOps
Leo Garciga, Director of Information Management, HQDA-DCS-G2, Army
Gus Coronel, Cloud Security Architect, Check Point
Payam Payandeh, Solutions Engineer, Datastax
Gavin Schmidt, Principal Engineer, MITRE

Challenge Area 4: Serverless Computing and the Impact on Cloud
Cameron Boozarjomehri, Principal Engineer, MITRE
Andrew Nebus, Senior Principal SME: Trusted Advisory, ASRC Federal

Challenge Area 5: Serverless Computing and the Impact on Cloud
Jeff Flick, Acting Director, Enterprise Network Program Office, NOAA
Sandeep Shilawat, ManTech Executive Director, Cloud/DevOps and IT
Rock Sabetto, Principal Systems Engineer, MITRE

Below is a list of government, academic and industry members who participated in these dialogue sessions:

Challenge Area 1: Cloud Smart & Impact on Cloud Adoption

Tan Luong, US Mint; Sam Lakhai, CDWG; Thomas Santucci, DOJ; Ben Todd, GSA; Audrey Payne, NRC; Wayne McFadden, Air Force; Lewis Quick, DHS; Noah DiDonato, Carahsoft; Sally Dillinger, Red Hat; Sonika Mohan, HHS; Luis Cano, Census; Terrance Glover, Green Electronics Council; Rahmira Rufus, MITRE; Tony Vicinelli, NLyte

Challenge Area 2: Challenges to Cloud Enabling Security
Michael Mather, DHS/CBP; Robert Fleming, GSA OIG; Donald Johnson, VA; Joseph Walter, MITRE; Lynette Wilcox, MITRE; Judith Hwang, Pyramid Systems; Shamil Hameed, UMD; Kimberly Hancher, ATARC Board

Challenge Area 3: Cloud-enabled Rapid Development and DevSecOps
Khalid Al-hassan, EPA; Gary Retzclaff, TRADOC G2; Jeff Diederiks, Pyramid Systems; Than Williams, Apptio; Jennifer Dostal, Peace Corps; David Hansen, MITRE; Vaqar Ahmed, USDA; Don Lamb, MITRE; Audrey Winston, MITRE

Challenge Area 4: Serverless Computing and the Impact on Cloud
Victor Pimentel, GSA; Christian Baer, NRC; Stephen Parowski, TTEC; Shane Cashman, Carahsoft; Ferhad Chohan, Red Hat

Challenge Area 5: Serverless Computing and the Impact on Cloud
Kevin Stern, CFPB; Chris Puccini, IBM; Pharist O'Neal, IBM; Ted Rutsch, ZScaler; Henry Davis, NRC; Jeff Flick, NOAA; Peter Morscheck, ManTech; Greg Braggs, FDIC

Thank you to everyone who contributed to the MITRE-ATARC Cloud & Infrastructure Collaboration Symposium. Without your knowledge and insight, this White Paper would not be possible.

Sincerely,

Tom Suder
Founder, Advanced Technology Academic Research Center (ATARC)

FEDERAL IT SUMMIT SERIES

JUNE 2019
FEDERAL CLOUD & INFRASTRUCTURE SUMMIT REPORT*

July 25, 2019

Justin F. Brunelle, Cameron Boozarjomehri, David Hansen, Christine Kim, R. Scott Paul, Quang Nguyen, Rock Sabetto, Gavin Schmidt, Mari Spina, Joseph Walter, Katy Warren, Adam Yee
The MITRE Corporation

Tom Suder
The Advanced Technology Academic Research Center

Contents

1 ABSTRACT

The most recent installment of the Federal Cloud & Infrastructure Summit, held on June 25, 2019, included five MITRE-ATARC (Advanced Technology Academic Research Center) Collaboration Sessions. These collaboration sessions allowed industry, academic, government, and MITRE representatives the opportunity to collaborate and discuss challenges the government faces in cloud computing and data center modernization. The goal of these sessions is to create a forum to exchange ideas and develop recommendations to further the adoption and advancement of cloud computing and data center management techniques and best practices within the government.

Participants representing government, industry, and academia addressed five challenge areas in the federal cloud and infrastructure domains.

- Cloud Smart & Emerging Policies – Impact on Government Cloud Adoption
- Challenges to Cloud Enabling Security
- Cloud-enabled Rapid Development and DevSecOps
- Serverless Computing and the Impact on Cloud
- Zero Trust and The Cloud

This white paper summarizes the discussions in the collaboration sessions and presents recommendations for government, academia, and industry while identifying intersecting points among challenge areas. The sessions identified actionable recommendations for the government, academia, and industry which are summarized below.

Emphasize and embrace agility when operating within and migrating to the cloud. This includes continuous learning and rapidly and safely achieving – and learning from – failure. Organizations should use small, incremental pilots or use cases to drive cloud activities.

Clarify accountability and authority. Particularly with technical staff and developers, the accountability of staff members when rapidly operating and delivering code is challenging. Organizations should clearly indicate staff member responsibilities and create interdisciplinary teams (e.g., with security personnel embedded with developers). Authorizing officials and authority to operate agreements should be appropriately balanced with agility and speed, especially as

the government adopts more modern software delivery practices within cloud environments.

Develop roadmaps for adopting cloud services. This will help provide concrete pathways for adopting emerging concepts such as Zero Trust and serverless computing. This will also help organizations set goals, identify requirements, and determine where tools or playbooks (e.g., cloud smart) are required.

Consider FedRAMP in the context of multiple cloud operation patterns. DevOps practices, security controls, and other cloud-based operations should leverage FedRAMP standards to ensure cloud security. However, many sessions cited that FedRAMP – while a good starting point for cloud security – must often be modified or supplemented to suit agency goals.

Treat cloud security as a primary goal. This will help emphasize the need for good data and cloud security hygiene, as well as guide teams toward adopting best practices for monitoring and appropriately leveraging resources available within the cloud.

2 INTRODUCTION

During the most recent Federal Cloud & Data Center Summit, held on June 25, 2019, five MITRE-ATARC (Advanced Technology Academic Research Center) Collaboration Sessions gave representatives of industry, academia, government, and MITRE the opportunity to discuss challenges the government faces in cloud computing and data center modernization. Experts who would not otherwise meet or interact used these sessions to identify challenges, best practices, recommendations, success stories, and requirements to advance the state of cloud computing and data center technologies and research in the government. Participants ranged from the CTO, CIO, and other executive levels from industry and government to practitioners from government, industry, and MITRE to researchers, students, and professors from academia.

The MITRE Corporation is a not-for-profit company that operates multiple Federally Funded Research and Development Centers (FFRDCs) [17]. ATARC is a non-profit organization that leverages academia to bridge between government and corporate participation in technology[1]. MITRE works in partnership with ATARC to host these collaborative sessions as part of the Federal Cloud & Infrastructure Summit. The invited collaboration session participants across government, industry, and academia worked together to address challenge areas in cloud computing and data center management, as well as identify courses of action to be taken to enable government and industry collaboration with academic institutions. Academic participants used the discussions as a way to help guide research efforts, curricula development, and to help produce graduates ready to join the work force and advance the state of cloud and data center research and work in the government.

This white paper is a summary of the results of the collaboration sessions and identifies suggestions and recommendations for government, industry, and academia while identifying cross-cutting issues among the challenge areas.

3 COLLABORATION SESSION OVERVIEW

Each of the five MITRE-ATARC collaboration sessions consisted of a focused and moderated discussion of current problems, gaps in work programs, potential solutions, and ways forward regarding a specific challenge area. At this summit, sessions addressed five topics.

- Cloud Smart & Emerging Policies – Impact on Government Cloud Adoption

[1] http://www.atarc.org

- Challenges to Cloud Enabling Security
- Cloud-enabled Rapid Development and DevSecOps
- Serverless Computing and the Impact on Cloud
- Zero Trust and The Cloud

This section outlines the goals, themes, and findings of each of the collaboration sessions.

3.1 Cloud Smart & Emerging Policies – Impact on Government Cloud Adoption

The Cloud Smart & Emerging Policies session focused on how Cloud Smart[2] has affected the process of government cloud adoption since the release of Cloud First in 2011 [11]. Cloud Smart is intended as a guideline for improving cloud adoption across all government agencies. Cloud Smart proposes three major accelerators to adopting cloud in federal agencies; this session discussed the practical implications of leveraging these accelerators, including the following topics:

- Identifying specific cloud security requirements
- Addressing the changes in acquisition and procurement related to cloud
- Acknowledging that cloud adoption drives changes in workforce skills and organization alignment and requires education for all staff

3.1.1 Challenges

The collaboration session discussions identified two primary challenges with incorporating the concepts of Cloud Smart to cloud adoption.

- **Changing skill sets and potential job loss:** There is still a fear that adopting cloud will eventually lead to a loss of jobs. Cloud brings with it an automation of processes which personnel may connect with the correlating job becoming obsolete.
- **Adapting to cloud:** Roles and government priorities will shift with the adoption of cloud. Agencies must begin to shift their mindset to adapt to the needs of cloud. There must also be a shift in mindset to working to be more agile for usage of cloud to be maximized.

[2] https://cloud.cio.gov/strategy/

- **Knowledge:** Cloud is still a relatively new technology for much of the Federal Government. Knowledge of how to integrate cloud capabilities into the government along with how to make and execute strategies and plans remains a challenge.

3.1.2 Discussion Summary

Discussions during this collaboration session focused on the impact that both Cloud First and Cloud Smart have had with government. Cloud Smart identifies three major accelerators to cloud adoption for the Federal Government.

The session started off with a discussion on challenges government organizations have faced with Cloud Smart adoption. Many agencies who have achieved success in cloud adoption have first faced significant failures. Those agencies have learned from their mistakes, regrouped and re-planned, and ultimately been able to achieve their IT modernization goals.

One significant and common challenge that many faced was cultural fear; there is a belief (usually not founded in available data) that jobs will disappear with the adoption of cloud. However, the session participants noted that the work and roles of people may shift with an additional need for extra education. The work need was still there – if not growing larger – due to the mission improvements and successes that cloud can bring an agency. Providing cloud education to all stakeholders can significantly reduce these concerns and at the same time, improve the ability to acquire, integrate, and use cloud services.

The discussion of workforce brought up another point in terms of cloud and cost surrounding it. Cloud is sometimes explored as a way to lower IT costs and reduce the amount of workforce needed. However, the participants cited that cloud impacts on IT costs are difficult to determine and often do not result in lower total costs. Instead, cloud often introduces opportunities for better mission capabilities, and usage-driven cost transparency and an abundance of available resources. An agency's costs while utilizing the cloud and its costs for an existing system are two very different values that are often not comparable.

There was also significant emphasis on the topic of security. Government cloud faces a unique challenge in that actions within the cloud can lose visibility due to TIC 3.0[3]. Although there are established ways to keep data secure, government cloud's strict policies can act as a deterrent to the overall security of a system.

Cloud Smart introduces opportunities for elastic TIC and Cloud Access Point (CAP) to remove potential bottlenecks. To ensure an agency's security is up to par, there was noteworthy discussion on creating pilot Use Cases to understand the needs of a specific agency and develop a security architecture capable of meeting them. These Use Cases could

[3] https://policy.cio.gov/tic-draft/

return measurable results that would aid in a creation of a risk assessment and determining bottlenecks within the cloud. Such outcomes with an overlay of Cloud Smart policies will allow for agencies to build their path forward in cloud adoption.

Further significant discussion involved cloud acquisition and procurement. Cloud Smart, while continuing to be dependent on federal acquisition regulations and policy, recognizes the need to adapt agency policy, processes and procedures to the flexibility the cloud offers. Agencies IT needs, and industry ability to meet them, change over time. Historically, cloud costs tend to lower over time, and new and improved services become more available and secure. Therefore, new and existing contracts must have the ability for agencies to takes advantage of these improvements in an agile and ongoing way.

Educating the cloud work force is key to successful cloud adoption and adaptation to cloud opportunities. The discussion revolved around providing knowledge of fundamental cloud concepts to everyone, including acquisition personnel, technical staff, security staff, end users, and agency leaders. The adoption of cloud will have impacts across the board, so the right training and education will be the fastest way to establish common understanding and knowledge for all stakeholders. Agile processes and cloud computing often go hand-in-hand. And a knowledgeable workforce is critical to ensuring good decision-making, planning a and execution of cloud in an ever changing environment. The work force and its culture can sometimes become the most underestimated problem during cloud adoption. Developing an appropriately skilled workforce with an iterative and agile mindset is key for successful cloud adoption.

The collaboration session ended with the discussion of where to start with cloud. Cloud adoption should not be viewed as a large jump or as an “all or nothing” decision. Small steps can be leveraged as learning opportunities for agencies and can build knowledge and confidence in tactical cloud migrations. These steps must also be guided by an overarching cloud adoption strategy that ensures the efforts of agency leadership, mission capability, acquisition office, IT office, and security create a synchronized execution of cloud adoption and adaption. Cloud Smart policy highlights areas to leverage as accelerators in this journey. Finally, operations in the cloud – including performance, security, and cost – should continually be reviewed and evaluated for potential improvements in efficiency and effectiveness.

3.1.3 Recommendations

The participants in the Cloud Smart & Emerging Policies collaboration session identified several important findings and recommendations.

Take steps towards cloud adoption: There is a heavily emphasized misconception that agencies must move all their infrastructure to cloud with one swift move. However, it was thoroughly discussed that cloud adoption can begin with small pilots and use cases to determine the best path forward to full (or nearly full) cloud adoption.

Educate everyone: Cloud adoption affects entire agencies. This requires everyone involved to understand the full processes and be educated on any differences in processes that will come with cloud. Once cloud has been put into full operation, it must continually and iteratively advanced and improved. Having everyone educated from the start will allow for an easier, streamlined work flow.

Utilize newly created models: Cloud Smart works as a strategy that guides government agencies into cloud adoption. Similarly, there are other models – most notably, Zero Trust – that have developed that agencies can leverage to begin incorporating cloud into their systems.

3.2 Challenges to Cloud Enabling Security

As industry migrates work load to the cloud, cyber security organizations struggle to maintain their highly-regarded Common Operating Picture (COP). But cloud-based security tools can be different from traditional tools in ways that may both enable or disable Standard Operating Procedure (SOP). The session participants cited that workload engineers are often not well versed in describing perceived threats to their environments and Security Operations Center (SOC) personnel are often not extensively experienced in cloud security tools. Consequently, the varied lexicon driven by the multitude of cloud offerings in the marketplace can cultivate coverage gaps in defensive posture after migration.

This break out session addressed the challenges faced by government organizations and their industry partners in maintaining or improving Defensive Cyber Operations (DCO) capabilities as the shift to cloud computing is made. Government leaders and industry professionals shared their experiences, challenges, and successes in this area. This collaboration session discussed the challenges, issues, and lessons learned of using cloud-based security tools.

3.2.1 Challenges

This session described – in depth – several pain points of government practitioners.

Lack of Control (Perceived or Otherwise): The largest pain point described by session participants to using cloud-based security tools was focused on obtaining sufficient visibility into an environment outside of their control. Even with the cyber-security tools provided by the CSP, there is a belief that the underlying CSP infrastructure (e.g., people, hardware) remains an undefined risk. While Service Level Agreements (SLAs) and FedRAMP [12] certifications attempt to mitigate the risk, the session participants indicated that the use of cloud services opened up their agency's perimeter to inherent risks that were not captured in CSP-provided cyber tools that focused on the VM/hypervisor layer and above. Due to the nature of FedRAMP and federal contracting constraints, the participants acknowledged that CSPs offer fewer choices in cybersecurity tools in their federal enclaves than within their generally available commercial environment.

Associated with this lack of security control when operating within a CSP, session participants worried about the management of CSP cyber-security auditing and monitoring tools – both those provided to CSP customers and those used internally by the CSP for operations. Agencies retain responsibility for safeguarding their data. However, in the case of a data spill, how can they guarantee that the data has truly been removed from the CSP's environment? Cloud-based encryption technology may someday be accredited to handle sensitive data; until then, moving any data into the cloud is considered risky, requiring an appropriate executive level of risk acceptance.

In the case of a data breach, agencies wondered if CSP internal security logs maintain an appropriate level of integrity to pass the 'chain of evidence' rules for digital forensics to be admissible in litigation. Conversely, are the agency's own cloud environment tool logs tightly controlled enough so that their own cloud-security tool logs would be admissible in court? Some agencies have attempted to mitigate this lack of control by creating deviations to FedRAMP requirements that have only made migration to the cloud more difficult.

Evolving Organizational Roles/Responsibilities: Federal cyber-operations teams do not have guidance or experience managing an external stakeholder, such as the CSP and the CSP's cybersecurity organization. How often should an agency perform "audits" to ensure that the CSP is proactive, not reactive, in responding and reporting on infrastructure risks outside of the control or visibility to the agency's internal cyber-operations team? This requirement for external stakeholder monitoring is not typically built into federal cyber-security operations contracts, leaving a gap.

With the introduction of cloud cyber-security tools, there is an increased need for agencies to provide training to their cyber-operations teams. These additional training costs are often not factored into the decision to migration to the cloud. In the past, agency cyber-

operations teams spent a great deal of time focusing on detection of related security events. With cloud cyber-security tools increasingly using automation and artificial intelligence to automatically correlate massive amounts of data, cyber-operations teams need training to focus on understanding the context of an event. Cyber-operations teams need to be able to correlate events happening in their cloud environment with events happening in the legacy agency environment.

New agency operational teams (e.g., acquisition, change management) need to be linked into the cybersecurity management process in order to take advantage of the increased speed and accuracy of new cloud-based cybersecurity analysis tools.

With an increased reliance on the shared responsibility model for security controls, agency business leaders need to move from managing security through audit/compliance checks to using a risk management approach towards security. A federal agency still retains responsibility for the role of cybersecurity in the cloud environment. Cloud-based cybersecurity tools need to feed agency risk assessment models for the agency to fulfill it's mandate to manage risk. However, the Risk Management Framework (RMF) [16] is relatively new for Federal agencies and will require training and resources for agencies to build their risk models and adopt a risk management approach to operating in the cloud.

Increased Complexity for Agency Cyber-Operation Teams: With the high number of controls provided through FedRAMP approved environments, it's hard for agency cyber-operations teams to determine who has operational responsibility to manage/maintain control. For example, session participants questioned what role(s) their cyber-operations teams and CSP cyber-operations teams had when meeting federal mandates for data encryption (e.g. FIPS 140 [14]). While continuous authorization remains a goal that will utilize cloud security tools, the current Authority to Operate (ATO) process for cloud applications takes longer because of this shared responsibility model.

Response planning requires tight integration between agency cyber-operations teams and the CSP. Clearly defined escalation paths need to be negotiated before an incident occurs. When outsourcing operations to a CSP, agencies had issues integrating ticketing systems, resulting in the concurrent use of multiple incident response systems.

Most agencies have adopted a policy of using multiple cloud providers, resulting in an exponential increase in operational complexity. CSPs will offer their own cybersecurity tools, but the agency cyber-operations team then has to do event management over multiple CSP cybersecurity tool stacks. The workload to maintain configuration management control across CSP environments also increases the workload on an agency's cyber-operations. The key to avoiding vendor lock-in for cyberoperations teams has been to keep some functions

outside of the CSP environment (e.g., don't use cloud cybersecurity tools).

The demand for cybersecurity professionals has added to the complexity of managing an agency cyber-operations team. Federal managers have to decide which tools/skills they must grow internally, and which they must outsource to CSP staff and cloud security tools.

3.2.2 Discussion Summary

As the move to the cloud continues, now with re-ignited thrust from the Cloud Smart initiative, government SOC leaders continue to ask, "What should we move to the cloud?" The answers continue to vary over time as the marketplace develops. To complicate the question, not all SOCs operate the same. Some focus heavily on incident response while others bring advanced threat detection systems to bear for event alerting. Depending upon the exposure of systems defended, some SOCs may be inundated with real-time incident response activities while others may spend more time on forensics and threat hunting. Data collection and data analytics are driving innovation in this area.

Understanding the volumes of data necessary to affect DCO and provide the effective COP can be a daunting task for most SOC leaders. Though estimation methods have surfaced over the last few years[4], SOC operators can easily see 1-2 terabytes of data per logging device over a year's time[5]. Multiply that by the number of event logging systems such as firewalls, network traffic analysis devices, and server hosts and the numbers scale quickly. Data retention cycles then become important and storage types for real-time, near real-time, and archive access drive procurements. The volumes of data and how to store them for usable access is only the tip of the iceberg.

SOC operators today are still wrestling with what to do with the data once they get it. At the same time, with volumes of data increasing and Security Information and Event Management (SIEM) systems improving, the cloud is delivering a paradigm shift from watching over the network to overseeing the use of cloud-based resources. Early threat detection systems leveraged network event logs and system interfaces. With the cloud, the end-point is getting increased attention. When the network is someone else's job to police and logs are only available for the systems that organizations deploy to the cloud or the services organizations consume in the cloud, cyber defenders quickly realize their inadequacy at understanding threats at the cloud edge. Moreover, with each CSP comes a different set of logging systems and event data types.

While the SOC struggles to understand the cloud edge, organizations are finding that

[4] https://www.linkedin.com/pulse/how-correctly-size-your-siem-investment-kerem-ozturk
[5] http://www.buzzcircuit.com/208/

the SOC is ill-equipped from both an education standpoint and a process perspective. Development teams have rarely handed over threat detection algorithms with their move to production. In fact, application engineers rarely consult with SOC operators about the types and forms of threats to look for in their application systems. As a result, not only do skills have to evolve but the processes for collaboration and knowledge transfer also have to evolve to make the enterprise effective at DCO in the cloud era.

It is generally perceived among system owners that the visibility available in the cloud is better than the visibility they had when in the data center. However, the data center always owned DCO and rarely ever made their COP available to system owners. So, while some of this perceived benefit is likely an artifact of the shift in responsibility from the data center management to mission owner management, it is clear that some CSPs are advancing the state of cloud service event logging and threat alerting. This is indeed providing an improved COP.

In addition to the cost of SIEM systems and associated network and storage compliments, getting data out of the cloud can add costs. Keeping cyber data generated by cloud-based systems in the cloud can be cost saving. To enable this, the CSP are leveraging artificial intelligence and machine learning platforms. Government practitioners are nearing the point at which cyber defenders need to have some background in data science. Because of skills gaps in these domains, adoption of these cloud-based DCO systems is not robust.

As government SOC leaders wrestle with the questions of how to "Cloud Enable" cybersecurity, industry and academia must understand their plight. It is one which is cost constrained and not rapid to convert. It involves skills and technologies that are constantly evolving just as fast – if not faster – than the threat environment. In the end, however, it is a business sector ripe for change and enthusiastic about the opportunities cloud-based DCO will bring.

3.2.3 Recommendations

The participants made several recommendations.

Replace on-prem cyber-capabilities with cloud-enabled cyber-capabilities: The question of what on-premise capabilities to move to the cloud was of high interest. Participants generally agreed that cyber visibility may improve when IT systems move to the cloud. The session participants also noted that keeping cyber event logs in the cloud could save money by avoiding network transfer fees. But most of all, participants appeared enthusiastic about the opportunity to leverage compute processing and scientific platforms in the cloud for improved threat detection.

While avoiding specific vendor solutions, session participants cited containers for detection (monitoring and auditing) tools among the most popular cloud cybersecurity tools. Cloud-based security tools for detection have been incorporation artificial intelligence and machine learning to significantly improve anomaly detection. There have been decreases in false positives, and cyber-operations teams are more quickly able to find root causes through artificial intelligence and machine learning-based correlation engines.

There was significant interest by session participants in the cloud-enabled method of implementing DHS' Continuous Diagnostics and Mitigation (CDM) [10] toolset, particularly at the Small Business Administration (SBA)[6]. SBA has piloted a "CDM in the Cloud" pilot in conjunction with DHS. The SBA pilot uses CDM tools deployed in the cloud (i.e., IaaS) versus the traditional DHS CDM implementation of tools into an agencies on-premises environment. The SBA had already cloud-enabled their TIC [9] cybersecurity capabilities through a successful pilot to move their TIC into the cloud[7].

Use application rationalization to speed up the ATO process: Session participants were achieving faster ATOs by using hardened baseline cloud-images of standard baselines that had already been vetted. Similarly, using accredited shared cloud-based security tools (e.g., Identity as a Service) is another method of speeding accreditation. Various cloud-based cybersecurity tools have been built into DevSecOps environments to achieve faster ATOs, as well.

NISTIR 8011, Automation Support for Security Control Assessments [8] provides a primer for agencies on how to use the CDM capabilities to achieve automated – and continuous – authorizations. The co-use of CDM and the RMF have the potential to transform how federal agencies perform assessment and authorization of their systems. Through the continuous monitoring suite of tools provided by CDM, agencies will be able to implement an on-going authorization process, phasing out their current authorization and accreditation programs (which typically go 3 years between re-accreditations). With data from cloud enabled CDM tools incorporated into their RMF programs, agencies will be equipped to make near-real-time determinations of their IT systems' security status, and use risk-informed, business-based rules for a course of action.

Virtualize Centralized Operations: Session participants expressed that their agencies have been using cloud-enabled cybersecurity tools to build a "single pane of glass" and cloud-based systems-of-systems to monitor their multi-cloud environments. This included cloud-based cyber tools to detect rogue systems within their environments, with automated

[6] https://fedtechmagazine.com/article/2019/05/sba-interior-energy-find-different-effective-ways-dep
[7] https://gcn.com/articles/2018/11/01/psi_transforming-sba-cybersecurity.aspx

quarantining of the system.

To incorporate new stakeholders into cyberoperations, agencies have been using the NIST National Initiative for Cybersecurity Education (NICE) Framework [15]. While originally intended as an education guide for workforce development, the NICE framework describes cybersecurity activities and roles that are independent of organizational boundaries (e.g., tasks to be accomplished not organizations to be staffed). As such, agencies can look beyond traditional in-house cyber-operations teams to find resources who can accomplish the necessary tasks presented in the NICE framework. This has proven helpful in building virtual centralized cyber-operations capabilities that incorporate external stakeholders, such as CSP operations teams, agency business leadership, and field operations teams.

Focus on data suitability and data architectures: Prior to settling on any cloud-based cybersecurity solutions, session participants agreed that it is important to first understand the agency's data architectures and data management requirements. These data structures drive the risk level and security profiles of an agency. Only have the data is understood appropriate cybersecurity solutions can be chosen, cloud-based or not.

Ensure fundamental cloud security is implemented: A key to successfully incorporating cloud-base cybersecurity tools is the application of basic cyber hygiene. Agencies must understand how well current cyber-capabilities are meeting basic cybersecurity functions, such as hardware/software inventory and identity. Perform an analysis of current strengths and weaknesses before looking to augment cyber-operations teams with cloud-based tools.

Make pro-active security a goal: With access to cloud-based big data capabilities, including artificial intelligence/machine learning, cyber-operations teams are learning to anticipate cyber-events. Linking cloud-based cybersecurity services can allow for automated detection or responses to cyber incidents –including recovery – in the cloud.

Be creative when looking for cyber skills: The NIST NICE Framework breaks down cybersecurity roles and actions in a manner that allows agencies to look for cyber-relevant skills in non-traditional locations within the organization. With the demand for cyber-professionals remaining very high, it is often easier to retrain staff with tangentially cyber-relevant skills. For example, internal auditors (data scientists) who monitor for waste, fraud, and abuse have transferable skills that can be applied to the cyber security field.

3.3 Cloud-enabled Rapid Development and DevSecOps

The Cloud-enabled Rapid Development and DevSecOps collaboration session discussed DevSecOps within the context of government cloud and infrastructure. The session discussed

the challenges faced by the government in adopting DevSecOps, as well as challenges with adopting DevSecOps in general. The discussion had an emphasis on how utilizing Infrastructure as Code (IaC) within DevSecOps poses unique challenges. The goals of this session were as follows:

- Discuss the modern software delivery process and the role of cloud
- Discuss recommendations or common challenges with using cloud in DevSecOps

3.3.1 Challenges

The collaboration session identified the following eight challenges with adopting DevSecOps within government infrastructure:

- Selecting the appropriate tools for your environment's DevSecOps toolchain
- Ensuring developers have good resources to help them take on the new challenge of managing infrastructure and security
- Identifying the security responsibilities of the developers while configuring the cloud infrastructure and understanding what security responsibilities are covered by the CSP
- Continuously monitoring the security of applications and the cloud infrastructure as tools, standards, and threats evolve
- Realizing there is a collaborative effort between automated testing, the RMF process, and FedRAMP that coordinate to achieve security
- Acquiring and maintaining the government workforce to execute DevSecOps
- Authorizing Official (AO) confidence in inheriting security from other certifications
- Giving developers accountability while recognizing that they operate within a large architecture

3.3.2 Discussion Summary

The session started with a brief review of security practices for cloud infrastructure. This included an emphasis on a Zero Trust model relying on deep packet inspection of all traffic on the cloud infrastructure and exercising rigorous configuration control to ensure the infrastructure complies with the selected standards.

The session then moved on to discuss challenges with implementing DevSecOps for government in the cloud. The first topic of discussion was how to select the right tools for the toolchain. A common challenge was combing through the large number of tools, especially since new tools come out all the time. The participants agreed that a place where tools can be compared, rated by users, and new tools can be reviewed would be useful and would improve the current practice of trial and error when it comes to tool selection.

Another challenge identified by session participants was understanding the additional steps the government needs to take to secure the DevSecOps infrastructure. Confusion often arises because many select their cloud products based on FedRAMP or other such accreditation which is supposed to provide some level of security assurances. Additionally, the CSP provides additional information on what parts of the Shared Responsibility Model for which they are responsible. With these differing information sources on security practices, it is difficult for the government to identify what additional measures they should implement, if any. A related problem is how to continuously monitor the infrastructure for security and ensure that it continues to meet security requirements once it is up and running. For example, the code may have been secure when the application was released, but if it has not been reviewed recently, it may no longer be secure. Developers need to be aware of not only their application but any dependencies it has.

Collaboration session participants then identified that achieving security in the DevSecOps environment is really a collaborative effort between multiple processes including FedRAMP accreditation, the RMF process, and automated testing. As part of this discussion, participants discussed a need to rework the RMF process to support DevSecOps. Additionally, participants identified a need to expand automated testing tools to include penetration testing. There was also discussion about how it will be difficult to automate penetration testing since it requires human ingenuity. The group also discussed that part of the value of DevSecOps is the ability to integrate continuous penetration testing into the design process since any detected vulnerabilities can be addressed in the next release.

The group recommended that government agencies formulate their requirements and communicate them to industry in order to drive services that better meet government needs.

The next major challenge discussed in the collaboration session was workforce. Participants identified a need for both contracting personnel and technical staff to have a technical and procedural understanding of DevSecOps. Contracting staff need to be able to procure the required environment and to understand what they are and are not getting from providers. Technical staff need to understand not just how to provision the tools and environments but what is available to them and what the contract allows them to do. Participants also identified

the need for security personnel that can verify the security assurances made by the CSP. A heterogeneous team of contracting, security, developer, and other technical personnel is necessary to effectively implementing DevSecOps practices.

The group also recognized that providing an Authority to Operate (ATO) to a DevSecOps application deployed in the cloud relies on the AO to be confident in the security controls that are being inherited from other certification processes. AOs need to be smart about what risk they can accept and recognize that they can always layer additional security on what is provided.

Another topic of conversation was using IaC as part of the DevSecOps process. Collaboration session participants recommended using a declarative rather than a scripted approach for deploying IaC to gain more control over the end result. The group recommended making use of configuration management tools that can review IaC definitions for compliance before they are deployed and monitor and correct the infrastructure after it is deployed.

Collaboration session participants also discussed tools that developers can use to check their code for vulnerabilities. One recommended approach was to check the Bill of Materials[8] for an application against the database of Common Vulnerabilities and Exposures (CVE)[9] to ensure there are no compromises in supporting code. Participants recognized this would not be effective for custom code and identified the Common Weaknesses Enumeration (CWE)[10] as a good reference for developers to check their own code.

A major challenge identified in the discussion is giving developers good guidance on what they should be looking for in the feedback loop of their DevSecOps process to ensure good security. Particularly when it comes to infrastructure, developers may not have the needed training and resources to identify problems they may be inadvertently creating. Collaboration session participants suggested making use of Information System Continuous Monitoring (ISCM) and Security Technical Implementation Guides (STIGs) as sources of information.

3.3.3 Recommendations

The participants in the Cloud-enabled Rapid Development and DevSecOps collaboration session identified several findings and recommendations.

Improve availability of tools: Establish a site for comparing and discovering DevSecOps

[8] A Bill of Materials is a manufacturing term referring to the list of materials needed to develop a product. In this case, the Bill of Materials refers to the tools, services, or other hardware and software components of a pipeline or DevSecOps environment.

[9] `https://cve.mitre.org/`

[10] `https://cwe.mitre.org/`

tools – perhaps including user ratings of the various tools. Formulate government requirements for DevSecOps tools and infrastructure and drive industry to meet them.

Adapt RMF for DevSecOps: Review and revise the RMF process to better support DevSecOps.

Automate penetration testing: Develop automated penetration testing tools to make approach practical for smaller-scale applications. Recognize that penetration testing requires human creativity and may be difficult to automate.

Transiton to Infrastructure as Code: Transition from a scripted approach to IaC to a declarative approach.

Leverage DevOps to deploy advanced security techniques: Take advantage of the unique features of DevSecOps to deploy advanced security techniques. Compare CVEs and CWEs to the Bill of Materials for an application before build to identify vulnerabilities. Use ISCM to monitor infrastructure. Utilize configuration management tools to monitor Infrastructure as Code definitions for security violations before deployment and to monitor/correct the infrastructure after deployment. Compare code check-ins to STIGs.

3.4 Serverless Computing and the Impact on Cloud

The Serverless Computing and the Impact on Cloud session focused on identifying critical elements necessary to transition any government application to the cloud, with specific considerations for moving directly towards a serverless computing architecture. As cloud solutions have advanced, cloud providers have developed solutions that shift the burden of server management off of the customer, enabling customers to build applications by simply writing code that can automatically run and scale on a platform without being burdened with managing their underlying system dependencies. Session participants expressed that this paradigm shift has made it difficult for agencies interested in adopting cloud to convey the benefits of serverless computing because the paradigm is considered too foreign for decision makers. This is especially true when those decision makers are more familiar with Containerization, Virtualization, and their existing legacy solutions all of which have a clearly defined place for a server within their architecture. As a result, session discussion centered on clarifying what exactly is serverless computing and how it addresses the needs of the provided use cases in terms of affordability, maintainability, interoperability, security, and

ease of adoption.

This session had the following four goals:

- Explore agency applications and use cases for adopting serverless computing
- Identify challenges and critical consideration for adoption
- Discuss how serverless computing compares to server centric models for the provided use cases
- Recommend guidance on what key criteria help an agency successfully transition to the cloud in addition to serverless computing

3.4.1 Challenges

The collaboration session discussions identified several challenges with adopting serverless computing.

- Serverless computing is difficult to explain to decision makers because the name "serverless" is confusing in terms of conveying proposed architectures and solutions.
- It is difficult for government to know if serverless computing is right for an agency if the agency does not fully understand its own business processes.
- If an agency does decide that serverless computing is the right choice, it is not clear what the roadmap to adoption looks like.
- It is difficult to convey the security and performance benefits of serverless computing because government prefers to have more control over the underlying resources and infrastructure rather than less.
- Agencies are weary of taking steps that they fear will lock them in to one solution provider who will not integrate well with other tools, technologies, or platforms.

3.4.2 Discussion Summary

The session began with introductions pertaining to the participants' organizations and critical use cases. This was needed given the frequent confusion surrounding the term; there was a fundamental need to better explain what exactly "serverless computing" was, especially since the name is a misnomer.

Simply put, serverless computing is a paradigm in which solution providers allow customers to simply provide code that they wish to run as processes. Despite the name, the provider is still creating servers that execute code provided by customers; however, they only do so on triggering events and carry the burden of ensuring any such server is appropriately provisioned, patched, and disposed of at the end of its use. The goal being that individuals using this paradigm will be able to spend more development time delivering application code that more closely delivers on core business process rather than spending time managing system dependencies. This shift also demanded a change in pricing models. Since code is not activated until there is a triggering event (such as hitting a RESTful endpoint), there is no need to pay for hosting or provisioning live servers. Customers simply pay for the length of time that the service is active in the routine execution of the provided code. This description was critical to helping drive discussion evaluating serverless computing in contrast to server-centric approaches.

This sparked further discussion around what it takes to develop and deploy on a serverless platform. One of the critical considerations was that developing code for a serverless platform demands efficient code that is atomic and can execute quickly. This concept was met with resistance as participants had concerns around the ease with which they might be able to update legacy code for execution on these platforms. This also led to discussion of the ease of integration with existing legacy systems that they had no intention of modernizing.

Other discussion centered on pricing and ease of portability between services. Despite serverless solutions being able to execute code in most popular languages, there was still concern regarding how easily one solution could be developed on one platform and transported to another. Furthermore, if performance is tied to speed of execution, then how does the language and typical user load affect an agency's ability to keep an application live? This allowed for a lively discussion of several inherent benefits of serverless computing.

- Since pricing is performance-based, it encourages efficient software development.
- Since there are no persisted servers, attackers find it more difficult to find and exploit vulnerabilities that are a result of out of date patches or system misconfiguration.
- Serverless code automatically scales with demand up to a customer defined threshold but also scales down, meaning customers don't pay for resources they don't use.
- Serverless tools offer superior analytics for user experience since every user's interactions follow routes that can be isolated and customized to specific user groups.

Other participants expressed considerations for performance in terms of applications. These concerns spoke to a need for identifying cases where data timeliness and system load needed to be addressed. Because serverless code is dormant until triggered, there is no value in using it for applications that infrequently need to communicate as quickly as possible. Similarly, if a service demands high graphical processing (e.g., for training machine learning models), serverless may not be the ideal resource because the cost to scale it adequately may be more expensive than a traditional server. That said, many human-centered use cases mentioned during the discussion seemed like ideal fits for serverless computing applications.

3.4.3 Recommendations

The participants in the Serverless Computing and the Impact on Cloud collaboration session identified several important findings and recommendations.

Clearly define terminology: Despite the potential benefits of serverless computing, it is clear it has "branding" issues where many are confused by the term "serverless computing". Participants felt this was a critical sticking point that cloud providers should address in order to see wider adoption.

Offer clear pricing: Despite the opportunity to be a price affective alternative to virtualization and containerization, the issue remains that cloud providers often make their pricing strategies fairly opaque. Though serverless computing is measured in 100ms [13] of processing time, this amount is difficult to understand and calculate at scale. Furthermore, session participants expressed that the mechanism by which providers settle on any price for any unit of computation feels arbitrary.

Clearly define business processes: In order for organizations to successfully transition to the cloud – and particularly to serverless computing – they first need a solid grasp on their core business processes. Participants recognized that serverless computing offered opportunities to more clearly link specific business processes to specific code paths. However, developing clean, effective code for a serverless environment demands a clear understanding of what business processes are critical to an organization's success. This was a key sticking point in that while the benefits were clear, some feared it difficult to justify the overhaul and re-evaluation of key business processes when their leadership felt there was no need since they already had a working legacy application.

Skip straight to serverless: The critical benefit of serverless computing is that agencies need not migrate servers and other provisioned elements from existing infrastructure. Developers can simply focus on writing application code that mimics existing business logic and business processes that can automatically execute in the serverless environment. Similarly, these services provide tools for transitioning data to the cloud to guarantee the best possible performance in conjunction with code.

Establish and support evangelists: Even with a clear understanding of serverless versus server centric architectures, participants expressed that buy-in for a transition to cloud needs to be exhibited at the developer level. Developers and management alike need to clearly understand and advocate that the benefits of switching to serverless means freeing critical staff from maintaining legacy software and hardware to focus on delivering code that mirror business processes and help the enterprise succeed. This also means investing in the time, training, and additional resources necessary for a successful transition to the cloud.

In summary this session involved an active exploration of serverless computing and its benefits. Participants were quick to engage and share use cases they were most eager to update as their agencies considered the transition toward cloud hosting. Still, the underlying issue remains: the main inhibitor to the adoption of serverless computing as a clear path forward is awareness.

3.5 Zero Trust and The Cloud

As part of IT modernization, enterprise network perimeters are changing. Government organizations are moving to the cloud and more users are connecting remotely via mobile devices. Traditional analogies between an organization's network perimeter and a moat surrounding a castle have become less relevant. Security approaches must adapt in order to protect the organization's information, which now resides across these increasingly amorphous networks.

Zero Trust[11] is a security concept predicated upon continuous authorization of all information exchanges between all combinations of users, devices, and applications. In other words, trust is established on a transaction by transaction basis in contrast to traditional models that grant ongoing access to the network (versus one application in a microsegmentation architecture) when the user initiates the session.

The intent of this discussion was to explore the concept of Zero Trust and how it applies to

[11] Zero Trust is a term created by Forrester researcher John Kindervag in 2010 [1]. Forrester has since published a series of reports on Zero Trust implementation and supporting technologies.

government organizations as they migrate resources to the cloud. This includes the following related goals:

- Describe the concept of Zero Trust
- Identify the primary components of a Zero Trust Architecture (ZTA)
- Discuss the challenge of transforming a legacy security architecture to ZTA
- Identify the state of the market in terms of commercial offerings and government adoption

3.5.1 Challenges

The session uncovered several overarching challenges with respect to Zero Trust and the cloud.

Understanding Zero Trust. A significant portion of the discussion involved helping the audience gain a basic understanding of Zero Trust principles and how they contrast with legacy security techniques and supporting technologies. These basics are covered in many publicly available documents, but there are some hurdles to clear for those new to the subject. First, a fair amount Zero Trust literature is written in theoretical terms, making it too abstract for less technical readers to relate to it. Second, industry literature is rife with conflated language intended to sell products. A concerted effort is required to analyze these materials in order to understand how they relate to a specific organization and its current security capabilities.

Architectural Complexity. Enterprise networks are now increasingly amorphous, consisting of remote users attempting to connect to a variety of cloud and on-premise environments. Creating a ZTA capable of supporting this new architecture with interfaces to all necessary components is highly complex.

People Resources. While the principles of Zero Trust are not new, they do require traditional security professionals to shift from a compliance assessment mindset to an inline authorization mindset. For some organizations, making this shift will not only require training in Zero Trust security technologies, but also a cultural shift from a reactive stance to a proactive stance.

3.5.2 Discussion Summary

This session focused on gaining an understanding of Zero Trust and discussing its benefits and challenges. Government leaders and industry professionals shared Zero Trust perspectives, offering insights to inform Zero Trust initiatives.

For organizations seeking to explore Zero Trust, one simple yet important concept is time relevance. Zero Trust is fundamentally about authorizing exchanges in line with business processes as they are being executed. This differs from many of the passive or monitoring capabilities typically exist in an organization's portfolio of security tools.

In order to create a ZTA, the organization must possess strong identity management and user device management capabilities, applications that interface with a Zero Trust control plane, and a Zero Trust control plane that is capable of authorizing and orchestrating interactions (amongst all user, device, and application combinations). Additionally, in a ZTA, authorization decisions can actively incorporate changing conditions such as user permissions, security policies, device locations, and emerging threats or vulnerabilities. The concept of a ZTA control plane with these capabilities can also be referred to as a "Trust Engine" [1]. This approach enables users to connect directly to resources (e.g., CSP services) without being routed through a centralized enforcement point such as an on-premise gateway, where remote users are required to establish a Virtual Private Network (VPN) connection. While this ability streamlines the architecture, Zero Trust principles do include the need for the organization to implement traffic visibility and monitoring capabilities. Supporting this Zero Trust principle across an amorphous topology is challenging and requires a significant architectural investment.

When Zero Trust is properly implemented and managed, security risk that can be greatly reduced as "least privilege" is continuously enforced. The organization's attack surface can be significantly reduced when compared to traditional network architectures that allow users to navigate laterally following authorization.

While ZTA adoption may seem like a daunting proposition, organizations can take comfort in the fact that they likely have identity and device management capabilities in place that can be leveraged as start points. Organizations may also have invested in segmenting resources for security and performance reasons, which is another step toward ZTA.

Environment complexity is also an important Zero Trust consideration. Many organizations are using multiple cloud environments in combination with legacy on premise environments. This hybrid situation increases complexity due to the uniqueness of each individual environment. To cope with this complexity, organizations are seeking techniques to federate identity and device solutions and apply security policies across their hybrid en-

terprise. Addressing environment complexity is larger than Zero Trust and extends to the organization's overall IT strategy and roadmap. New and updated capability offerings to address these challenges are being brought to market by Cloud Access Security Brokers (CASBs), CSPs, traditional security vendors, and others. Market leaders are investing engineering resources to align new capabilities with the Zero Trust principles. Given the dynamic and complex nature of organizations' enterprises and the marketplace, creating an actionable ZTA roadmap is essential for success.

3.5.3 Recommendations

Session participants identified several important findings and recommendations.

Consider ZTA Pilots: Organizations should consider piloting select Zero Trust capabilities prior to deploying enterprise wide.

Create ZTA Roadmaps: Organizations should create clear roadmaps from current state to future state. These roadmaps should be part of, or synchronized with, cloud adoption roadmaps. ZTA Roadmaps should address application rationalization, as the ability to support Zero Trust may impact application decisions (e.g., may change required level of refactoring). The ZTA Roadmaps should also contain security tool rationalization. There are arguably multiple capabilities that are not primary Zero Trust components that will need to be kept in the overall security architecture. Examples include SIEM tools and vulnerability scanners.

Prepare for a learning journey: Organizations should be prepared to embark upon a Zero Trust continuous learning journey. Engage with NIST and DHS to help shape Zero Trust pilots, gather lessons learned from others, and stay abreast of emerging policies and standards. Explore what industry is creating and be prepared to analyze new offerings. Industry is rapidly responding to demand signals, but care must be taken to discern between engineering capability and marketing hype.

4 SUMMIT RECOMMENDATIONS

The June 2019 Federal Cloud & Infrastructure Summit collaboration sessions revisited some of the traditional recommendations (e.g., the emphasis on cloud security) well as identified new recommendations based on emerging concepts and technologies. As cloud becomes

more ubiquitous, the recommendations for overcoming perennial and emerging challenges are becoming more specific and tactical.

The government has long understood – yet struggles to embrace – the need to be agile when migrating and operating in cloud environments. Enabling and learning from failure is an understood necessity, but creating an environment in which this is possible can be its own challenge. The summit participants – once again – emphasized the need for small, rapid pilots to refine an agency's method of operating within a cloud.

Along the same lines, the participants often recommended adopting and leveraging emerging policies such as Cloud Smart. These policies are in response to the prior calls for government guidance on cloud migration. Developing organizational roadmaps will also help identify roadblocks, requirements, tools, and other aspects of the migration/adoption necessities. This is particularly important for government organizations that are adopting emerging concepts such as Zero Trust, serverless computing, or updating the way in which software is developed (e.g., via DevSecOps practices). The session participants also cited existing standards – such as FedRAMP – in the context of new cloud operating models. Often, the participants cited FedRAMP as being a starting point for addressing cloud security, but agencies must often adapt the policies to suit specific needs.

Security has been primary topic of discussion at past summits [7, 6, 5, 3, 2, 4] and this year's summit was no different. However, the emphasis this year was on new policies (e.g., TIC 3.0) and the relevant authorities and accountability during operation within the cloud. The participants recommended clearly defining accountability and authority with increasingly complex AO and ATO relationships. This is particularly important for DevSecOps since developers are often responsible for rapidly developing, testing (e.g., for security considerations), and deploying code within a cloud environment.

5 CONCLUSIONS

As in past summits, the June 2019 Federal Cloud & Infrastructure Summit collaboration sessions provided participants an opportunity to discuss the challenges, best practices, and recommendations regarding government use of cloud computing.

A prominent theme throughout the collaboration sessions was that cloud is becoming more ubiquitous. It is no longer a question of "how will my organization use the cloud?" The collaboration discussions have evolved to "my organization is using the cloud; how will my organization evolve its practices as a result?" This evolution is evident with the discussion focusing on cloud as an enabling tool rather than a conceptual end state that government

organizations need to achieve.

As always, security and government policy were focal points with TIC 3.0, Cloud Smart, and ZTA being major points of emphasis that cut across sessions. DevSecOps and the way in which developers and other practitioners use the cloud – with appropriate authority and accountability – received a similar emphasis. The tools that are used in DevSecOps (e.g., the development and security monitoring stacks) was an interesting topic. The participants in the session noted that tools were an enabler, which is an indication that the government is beginning to understand DevOps as a practiced discipline rather than a toolset that can be acquired.

Cloud is becoming more ubiquitous; its role in government efforts to improve efficiency demonstrates that cloud services are increasingly used, valued, and understood. This also creates new strains on the workforce, creating an ever-present lack of effectively trained and maintained government staff knowledgeable in cloud and cloud-based development practices.

ACKNOWLEDGMENTS

The authors of this paper would like to thank The Advanced Technology Academic Research Center and The MITRE Corporation for their support and organization of the summit.

The authors would also like to thank the session leads and participants that helped make the collaborations and discussions possible. A full participant list is maintained and published by ATARC[12].

REFERENCES

[1] ACT-IAC. Zero Trust Cybersecurity: Current Trends. Technical report, ACT-IAC, 2019.

[2] J. F. Brunelle, S. Anand, G. Barmine, M. Spina, K. Warren, A. Winston, M. Javid, A. Kemmer, C. Kim, S. Masoud, T. Harvey, and T. Suder. August 2017 ATARC Federal Cloud & Data

[12] https://atarc.org/event/cloud-infrastructure-summit/

Center Summit Report. Technical report, The MITRE Corporation; The Advanced Technology Academic Research Center, 2017.

[3] J. F. Brunelle, S. Anand, R. Cagle, C. Kim, M. Kristan, M. Spina, K. Warren, T. Harvey, and T. Suder. February 2017 ATARC Federal Cloud & Data Center Summit Report. Technical report, The MITRE Corporation; The Advanced Technology Academic Research Center, 2017.

[4] J. F. Brunelle, A. Bognar, V. Dhawan, N. G. Parrish, A. King, V. Kuppusamy, M. Malayanur, T. Harvey, and T. Suder. June 2018 Federal Cloud & Data Center Summit Report. Technical report, The MITRE Corporation; The Advanced Technology Academic Research Center, 2018.

[5] J. F. Brunelle, D. Davis, N. Gong, D. Huynh, M. Kristan, M. Malayanur, T. Harvey, and T. Suder. July 2016 ATARC Federal Cloud & Data Center Summit Report. Technical report, The MITRE Corporation; The Advanced Technology Academic Research Center, 2016.

[6] J. F. Brunelle, D. Davis, D. Huynh, M. Malayanur, B. Natale, H. Small, T. Harvey, and T. Suder. January 2016 ATARC Federal Cloud Computing Summit Report. Technical report, The MITRE Corporation; The Advanced Technology Academic Research Center, 2016.

[7] K. Caraway, N. Gong, J. Packer, J. Vann, J. F. Brunelle, T. Harvey, and T. Suder. July 2015 ATARC Federal Cloud Computing Summit Report. Technical report, The MITRE Corporation; The Advanced Technology Academic Research Center, 2015.

[8] K. Dempsey, N. Goren, P. Eavy, and G. Moore. Automation Support for Security Control Assessments: Software Asset Management. Technical Report NISTIR 8011 Vol. 3, National Institute of Standards and Technology, 2018.

[9] Department of Homeland Security. Trusted Internet Connections. `http://www.dhs.gov/trusted-internet-connections`, 2016.

[10] Department of Homeland Security. Continuous Diagnostics and Mitigation. `https://www.dhs.gov/cisa/cdm`, 2019.

[11] Federal CIO. Federal Cloud Computing Strategy. `https://cio.gov/wp-content/uploads/downloads/2012/09/Federal-Cloud-Computing-Strategy.pdf`, 2011.

[12] FedRAMP PMO. FedRAMP. `https://www.fedramp.gov/`, 2015.

[13] R. Gancarz. The economics of serverless computing: A real-world test. https://techbeacon.com/enterprise-it/economics-serverless-computing-real-world-test, 2019.

[14] NIST. Security Requirements for Cryptographic Modules. Technical Report FIPS PUB 140-2, National Institute of Standards and Technology, 1994.

[15] NIST. National Initiative for Cybersecurity Education (NICE) Cybersecurity Workforce Framework. Technical Report Special Publication 800-181, National Institute of Standards and Technology, 2017.

[16] NIST. Risk Management. https://csrc.nist.gov/projects/risk-management/risk-management-framework-(rmf)-overview, 2019.

[17] The MITRE Corporation. FFRDCs – A Primer. http://www.mitre.org/sites/default/files/publications/ffrdc-primer-april-2015.pdf, 2015.