BREAKING DEFENSE PATHFINDER

Building Multi-Cloud in the Intelligence Community

Information superiority is the hallmark of multi-cloud environments, along with automation, cybersecurity, and governance.

A Breaking Defense special report sponsored by Booz | Allen | Hamilton®

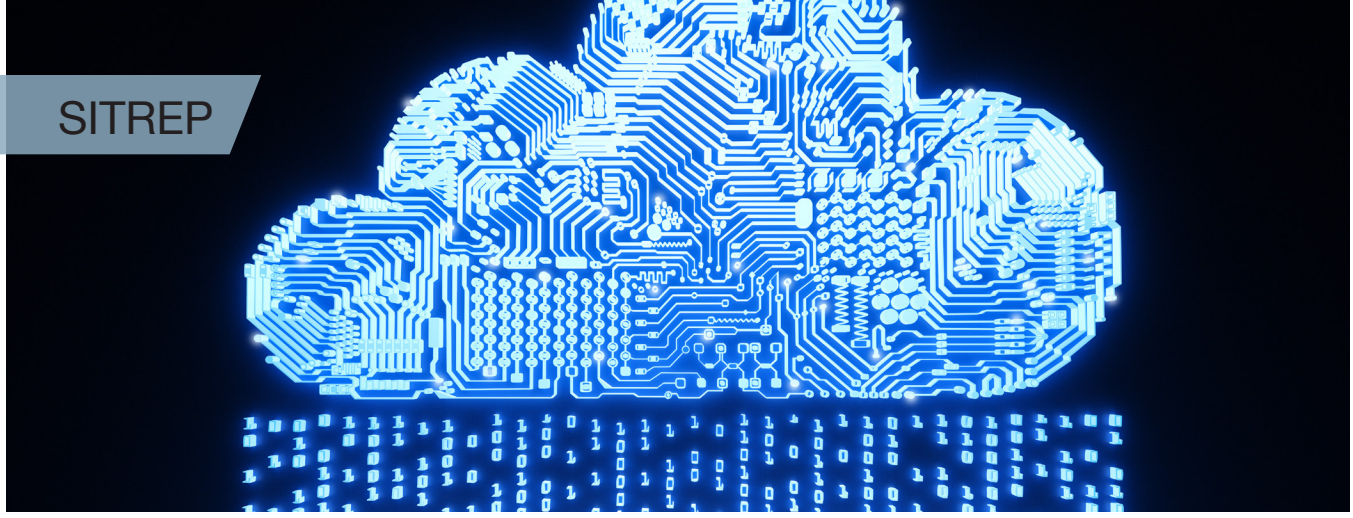

How Multi-Cloud Can Be A Force Multiplier for the Intelligence Community

The Intelligence Community (IC) has emphasized that the second epoch of its information technology enterprise will be defined, in part, by the creation of an integrated, vendor-flexible, interoperable, and secure cloud ecosystem. This environment will need to provide data collectors and analysts across the IC with access to the latest tools and technologies such as artificial intelligence (AI) and machine learning (ML), and must support workflows within and across multiple security domains.

The CIA's Cloud Integration and Multi-Cloud Management (CIMM) contract will serve as a lead effort for all IC multi-cloud efforts.

In such an environment, applications, services, data, and capabilities are best supported by multiple cloud environments—including private clouds and public ones from commercial Cloud Service Providers (CSPs)—as well as Software as a Service (SaaS) solutions. It's clear that the IC's future-state IT landscape will be a multi-network, multi-domain, and multi-cloud environment. However, the vision of exactly what that end state will look like, the best approaches to achieve it, and the means of measuring success are still maturing.

Universally agreed upon, however, is that the end goal of creating a multi-cloud environment is to gain information superiority and speed to mission through advanced storage, compute, and analytic capabilities not available in existing on-premise operations.

On-prem is not going away any time soon because of the vast infrastructure of existing unclassified, secret, and top-secret sensitive compartmentalized information (TS/SCI) networks, not to forget the Five Eyes coalition environment. What multi-cloud brings to the equation beyond that is flexibility and the ability to maximize open-source and publicly available information as a core component that is additive to what the IC is generating. CSPs also bring to bear a set of valuable, managed services not necessarily available on-prem—especially in the areas of automation, cybersecurity, and governance.

Multiple clouds also means multiple vendors, preventing the IC from becoming locked in to any one company and its proprietary products. It also exposes the importance of proper oversight and 24/7/365 management.

The IC is already preparing for a multi-cloud eventuality and the need for oversight through a solicitation for the CIA's Cloud Integration and Multi-Cloud Management (CIMM) contract, which will serve as a lead effort for all IC-agency multi-cloud efforts, including those at the National Reconnaissance Office, the National Geospatial-Intelligence Agency, the Office of the Director of National Intelligence, and others like the Defense Threat Reduction Agency (DTRA) and the Defense Intelligence Agency. As its name states, CIMM will help the IC stitch together its myriad clouds, act as the go-between between the government and CSPs, and provide governance and training. //

– Barry Rosenberg, Contributing Editor, Technology & Special Projects

The State of Multi-Cloud in the IC

IC organizations are continuing on their journey transitioning to the cloud, but agencies are in different phases of implementation. Some are already using cloud services to house systems, support and/or develop production workloads, and deploy SaaS applications like DoD365. Other agencies are starting to explore some of the feasibilities of the cloud, answering questions about how it will improve mission execution.

Their position on the yardstick depends upon funding, mission drivers, and the strategies outlined by each individual agency's CIO.

"The majority of the community, whether it's DoD or the IC, recognize the benefits of the cloud," said Christopher Brazier, chief technology officer and mission IT infrastructure and services department chief of DTRA's Information Management and Technology Directorate.

'You have managed services, and get scalability, elasticity, and flexibility with a cloud service offering.'

— Christopher Brazier, Defense Threat Reduction Agency

"It's the opportunity to innovate without procurement-related delays. Once you get that cloud-services contract in place, you have the ability to pay as you go using an on-demand model. You can set aside reserved instances for consistent workloads.

"Basically you're looking at increased cost efficiencies across the board. You have managed services, and get scalability, elasticity, and flexibility with a cloud service offering. The DoD and the IC recognize the need to have architectures in place to support multi-cloud, given some of the drivers that we're seeing today such as cost savings, resiliency, portability in supporting multiple cloud-service providers, and improved enterprise cloud management.

Brazier says that many federal organizations are in what he calls the "rationalization phase" where they've begun proof-of-concept pilot programs that include multiple elements such as inventorying applications and services to determine which ones to re-factor, rehost, replatform, or retire.

Transition Planning is the Key

Defining a blueprint for transition to cloud architecture is particularly important. For example, federal agencies moving from legacy networks such as the General Services Administration (GSA) Networx contract to GSA's newer and more secure Enterprise Information Systems network contract always insist that contractors bidding for the job include a detailed transition plan that ensures no interruption in agency services.

Transition leads to a full "realization phase" where agencies can build a greenfield architecture

for born-in-the-cloud applications where they can take advantage of SaaS and platform-as-s-service (PaaS).

Brazier points to a number of IC organizations that have already begun transitioning to this realization phase, particularly the Air Force's Platform One DevSecOps enterprise services team (which consists of software developers from it's Kessel Run and Kobayashi Maru software factories). Platform One provides DevSecOps managed services with collaboration and cybersecurity tools, source-code and artifact repositories, and DevSecOps as a Service so that development teams can simply employ those tools and pay per use at scale with bulk licenses.

For its part, Brazier says that DTRA is in the "limited-scope migration phase" where it's using a hybrid, public cloud architecture to support a DevSecOps pipeline already in production today.

"The majority of our production workloads and data are resident on-prem in our private cloud, and as we continue to progress in our journey towards the cloud my goal is to migrate our production on-prem workloads to the cloud," said Brazier. "Full realization for us means using cloud-native resources where it's appropriate. That would put us in a position to adopt new services and features much faster than before, which further helps accelerate delivery of capabilities to our warfighter/end user."

Many federal organizations are in what Brazier calls the 'rationalization phase' where they've begun proof-of-concept pilot programs.

DTRA is also focused on portability and accessibility with application workloads running on-prem in the agency's private cloud and in one or more public clouds.

"The key is staying in lockstep with the fast pace of technology evolution and the service-catalog updates that the cloud service providers put out," explained Brazier. "Our portability comes from containers, so both front-end and back-end applications are containerized. Added to that are features like Kubernetes or advanced container orchestration plus open application programming interfaces and standards. With the accessibility to directly deploy and scale services, we are able to meet growing and emerging needs and threats that we need to pivot toward and address on a daily basis."

For this effort, Booz Allen is DTRA's partner in modernization. The work involves expanding the agency's service offerings in the cloud by applying machine learning and artificial intelligence to relevant data to provide DTRA with operationalized data that the agency can use to improve the execution of its mission set.

"What the cloud is already providing for our development, test, and integration needs—the elasticity, flexibility, scalability—is a force multiplier," said Brazier. "It will help the broader intelligence community provide our warfighters and analysts with relevant capabilities, applications and tool sets, and permit DTRA to align to our agency's mission, which is to detect, deter, and defeat our enemies, as well as support our division motto of mission-driven innovation."

Motivations Driving the IC to the Cloud

As Brazier mentioned, portability and accessibility are two of several, driving motivations for the IC to pursue multi-cloud solutions and capabilities. Because these make up the business case for cloud expenditures, it's worth delving into them to a greater extent.

"Developing a set of ubiquitous platforms that allow applications to be deployed seamlessly in any compliant environment provides a legitimate

mission advantage," said Melissa Sutherland, Booz Allen's senior vice president supporting defense intelligence and modernization across the IC. "Portability is especially important for agencies in the IC and DoD that have compelling mission needs that may require them to deploy a set of applications into a robust set of platforms and infrastructures that can accommodate varying levels of connectivity, specific in-theater needs, or specialized equipment."

'As the technology, tools, and services related to multi-cloud continue to evolve, organizations will be able to adopt aspects of the multi-cloud promise more easily.'

— Melissa Sutherland, Booz Allen Hamilton

Multi-cloud architectures are also inherently more resilient, as they can achieve very high availability even through a CSP outage by leveraging multiple zones and regions that are more cost and performance effective.

Flexibility and access to innovation that can be spread around the IC agencies are additional reasons for transition to multi-cloud.

"The breadth of capabilities offered across the major CSP players and the fierce competition among them to differentiate themselves makes having broad access to more than one provider services a real advantage in tools and environments," said Sutherland. "Maintaining a range of cloud environments, cloud-native services, and capabilities helps keep the IT workforce productive and reduces the transition burden of implementing new technologies."

Improved cloud management is another area that IC agencies can benefit from. An enterprise multi-cloud approach to integrated management provides an opportunity to increase transparency, reduce risks, and increase efficiencies. Integrated governance can also help to ensure that application-level and CSP-level decisions are in alignment with the agency's strategic enterprise goals and objectives.

Finally there's affordability. Cloud management organizations can track the ever-changing cost structure and incentives of various CSPs and take advantage of the best pricing. This will also increase competition among vendors to drive costs even lower.

"As the technology, tools, and services related to multi-cloud continue to evolve, organizations will be able to adopt aspects of the multi-cloud promise more easily and in a sequence that best aligns with their broader business and mission objectives and strategies," said Sutherland.

Defining Multi-Cloud Today, and Tomorrow

Multi-cloud is still something of an evolving concept but it can be defined at a basic level as the concurrent use of multiple public and private clouds to meet an organization's mission, business, and IT requirements.

The term is also being used to describe architectural models that prepare an organization for the likely future evolution of cloud technology. This

includes several strategies that agencies are embarking on now or plan to implement in the future, include the following:

ENTERPRISE LEVEL–INTEGRATED MANAGEMENT: This view of a multi-cloud enterprise is a natural evolution of multiple vendor-proprietary CSP infrastructure, platform and software implementations and legacy on-premises capabilities. In this strategy, the enterprise is managed through a combination of centralized governance, policy and standards, cross-platform integrated services utilizing standardized application program interfaces (APIs), and a selection of automated tools.

Examples of central management services may include tools to support governance, audit, identity and access management, classification management tools, common service catalogs or storefronts, cybersecurity posture, platform health, and financial/cost accounting services. In this model, there may also be common management services specific to the CSP (e.g., encryption, firewall, event detection).

The speed, variety, and sheer volume of incoming data requires carefully chosen tools for data extraction, correlation, and enrichment.

Each application is still only deployed to a single cloud environment, but the underlying infrastructure support services may be provided at an enterprise level. An example of this type of solution is Microsoft Azure Arc, which provides centralized cloud management capabilities across any infrastructure environment, with an emphasis on managing container-based applications. Other SaaS and PaaS providers like IBM/RedHat and ServiceNow, also provide third-party solutions for multi-cloud management.

ENTERPRISE LEVEL–APPLICATION ANYWHERE: Once integrated management is established, organizations can start to leverage architectural principles, enabled heavily by Kubernetes clusters, across multiple cloud environments that allow container-based services to be uniformly deployed to any platform, providing maximum flexibility and portability. Under this model, an application can be deployed to a single cloud environment or to multiple environments. One solution to achieve this is Google Anthos, a managed Kubernetes services platform that allows containerized applications to be deployed across Google Cloud Platform, AWS, or on-premise infrastructures.

SYSTEM LEVEL–DISTRIBUTED APPLICATIONS: The next evolution of multi-cloud is for each application to consist of a collection of sub-services that can be distributed across multiple cloud environments. These sub-services can be architected to be cloud agnostic, or to take advantage of specific native cloud services via integrations (APIs).

Under this model, an application may be running in multiple clouds to support the system at the transaction level. This is the most complex and most nascent definition of "multi-cloud" in use today, but is perhaps the purest concept of the projected future of multi-cloud.

Another emerging trend that may enable distributed applications is the rise of PaaS and SaaS solutions to support integrated API management across a distributed ecosystem of cloud landing zones.

Applying Cloud Tools to Open-Source Intelligence

The architectural models defined above might vary somewhat over time, but they are far from aspirational as the amount of intelligence being collected is increasing exponentially. The speed, variety, and sheer volume of incoming data requires carefully chosen tools for data extraction, correlation, and enrichment—specifically tools that aren't nec-

essarily resident in agency on-prem operations.

"The more opportunity that the IC has to capitalize on native services that the CSPs are providing, the better they can execute their missions because so much data is being produced that they cannot build what they need to operationalize it in-house; it doesn't make sense to use custom code," said Dan Tucker, vice president of cloud and digital solutions within Booz Allen's Strategic Innovation Group, noting that that the amount of data available to IC agencies doubles every year.

'Think about the amount of data that is increasing exponentially. especially the data that's going to be needed to train these algorithms and these capabilities.'

— Dan Tucker, Booz Allen Hamilton

"With the onset and desire to use machine learning and AI-based applications that are going to ride in these multi-cloud environments, think about the amount of data that is increasing exponentially, especially the data that's going to be needed to train these algorithms and these capabilities; that's going to be very important to the intelligence's mission," said Tucker.

In particular, the IC is leaning toward pulling in publicly available information such as social media postings, known in the intelligence community as OSINT (open-source intelligence), that's available in the unclassified domain. Valuable OSINT that's identified by machine learning is then coupled with classified data to provide a holistic threat view. The tool set to perform such analysis is readily available in a multi-cloud environment.

"The way they categorize this or describe it is tradecraft automation," said Sutherland. "Analysts today are spending 80 percent of their time working in stovepiped environments with niche tool sets trying to collate that data in whatever environment they have. Whereas automating the environment with cloud-native tools lets them focus 80 percent of their time on analysis and 20 percent on analyzing the data."

Multi-Cloud to Execute IC Missions

The IC would be the first to agree with Tucker's assessment that the quantity of data available to them is growing exponentially. Finding value in that increasingly large data lake at the same rate at which it is collected virtually requires the AI/ML tools that work most effectively in the cloud.

"For us, to detect the threats, cloud architecture enables processing of complex models," said Brazier. "There's no limit, or virtually no limit, to data volume. Only budget constraints, right?"

For example, DTRA's visualization, modeling, and simulation group designs 3D models to characterize and simulate the scenarios of the nation's adversaries and provide its mission partners with models to make informed decisions. Deterrence is another one of DTRA's missions, and cloud enables the modeling and simulation group to be more forward leaning in terms of innovation and collaboration.

"The drive toward innovating real mission solutions depends on increased footprint to inject and validate new technologies," said Brazier. "Our adversaries are modernizing their military operations with advanced technology and artificial intelligence concepts, and we need to leverage the cloud to stay ahead of that. Our preferred way to stay ahead of our adversaries is by focusing on the development and delivery of those solutions."

That means the traditional acquisition approach to acquiring hardware and software necessary for on-prem operations to stay forward of the threat is no longer an acceptable first option. In fact, it's become more of a mission inhibitor than a mission enabler.

Conversely, cloud data and logic can be accessed from anywhere at any time once the requester or consumer is authenticated and authorized. Analysts can collaborate across multiple areas of responsibility such as Europe, Africa, or Asia, and build upon each other's concepts to deter the adversaries.

Multi-cloud architectures also enhance deterrence through resilience made possible by built-in high availability and elasticity. For example, multi cloud offers baked-in components for failover that can rapidly provision compute resources to continue to perform complex ML-enabled computations.

"Cloud-service providers focus their resources and efforts on guaranteeing uptime, which is re-

ally important," said Brazier. "So rapid response and the speed to develop and deliver new capabilities to our users down range is vital. There is no time to procure the hardware, configure the virtual machines, and focus on rapid development to directly respond to the emerging mission threat."

That concern is particularly germane to the Great Power competition against China and Russia, as well as other countries with advanced air defenses like Iran and North Korea. In that contest, speed to mission is vital to deterrence that prevents potential war.

The traditional acquisition approach to acquiring hardware and software necessary for on-prem operations to stay forward of the threat is no longer an acceptable first option. In fact, it's become more of a mission inhibitor than a mission enabler.

Arguably, one of the most important functions of multi-cloud architecture is to facilitate continuous integration and continuous deployment using DevSecOps processes that were discussed earlier. This creates and accelerates a path to production for getting an ATO (Authority To Operate) for an accredited application that can pull in data in and make sense of it in a cloud environment.

"Whether it pertains to the Great Power competition or a strategic intelligence question, analysts usually have a set of information needs that they're applying against the datasets," explained Sutherland. "Quickly developed applications using DevSecOps in the cloud will let them address those needs with intelligence products that create usable intel that is traceable and defendable back to the data sources."

A Roadmap for the IC Multi-Cloud

To achieve a multi-cloud state, Booz Allen suggests that the IC develop a strong foundation of cloud management maturity and implement open, well-architected, container-based platforms that will allow multi-cloud capabilities to evolve based on business needs and mission imperatives. The following represents potential evolutions of multi-cloud that IC organizations should plan to achieve:

INTEGRATED MANAGEMENT: Implement a central management service to include tools that support governance, service catalog management, and cybersecurity posture. There needs to be some aggregation of capabilities and performance to track multiple providers.

APPLICATION ANYWHERE: Leverage architectural principles across multiple cloud environments that allow container-based services to be uniformly deployed to any platform, providing maximum deployment flexibility and portability.

DISTRIBUTED APPLICATIONS: Enable a single application to consist of a collection of sub-services that are distributed across multiple cloud environments. These sub-services could be architected to be cloud agnostic or take advantage of native cloud services via integrations (APIs).

At present, architecting cloud-agnostic solutions may require an unreasonable reliance on third-party solutions and may increase complexity and reduce potential benefits from using cloud native services and the ongoing innovation that comes with them. But making some deliberate, smart bets on cloud-native can provide the best balance of innovation, cost effectiveness, and performance. IT leaders should continue to rely heavily on the transition to cloud-native applications and services to the highest extent possible, while managing the level of effort and cost associated with replacing services or moving to third-party services that provide more portability.

Organizations should develop and grow a taxonomy and architecture of common cloud services that will be implemented at a centralized enterprise level across multiple clouds, and those that will be implemented in common at a local cloud environment. As an organization's cloud platform grows and the offerings mature, the mix of central versus localized common services may change over time. //